(LEARNED) PHYSICS-BASED HUMAN MOVEMENT: SHARED MODELS FOR ANIMATION, ROBOTICS, VISION, AND BIOMECHANICS

Michiel van de Panne

Dept of Computer Science
University of British Columbia

“The best way to predict the future is to create it.”

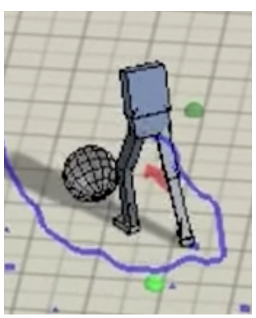
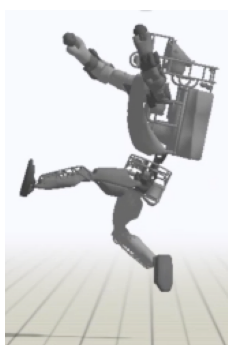



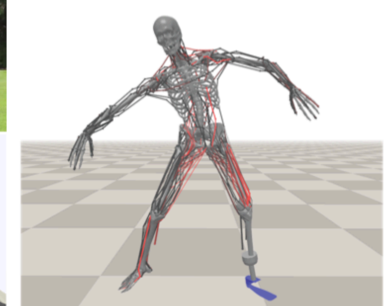

Dynamic Motion Control

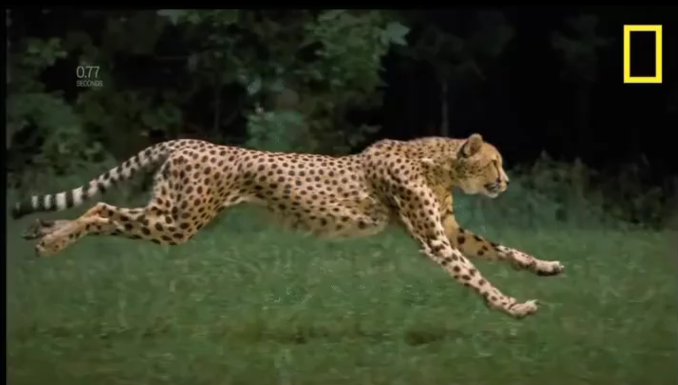

DEEP RL RESULTS

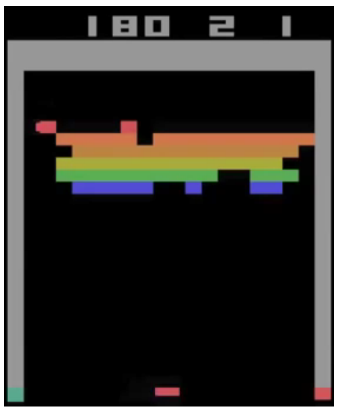

[Mnih et al. 2015]

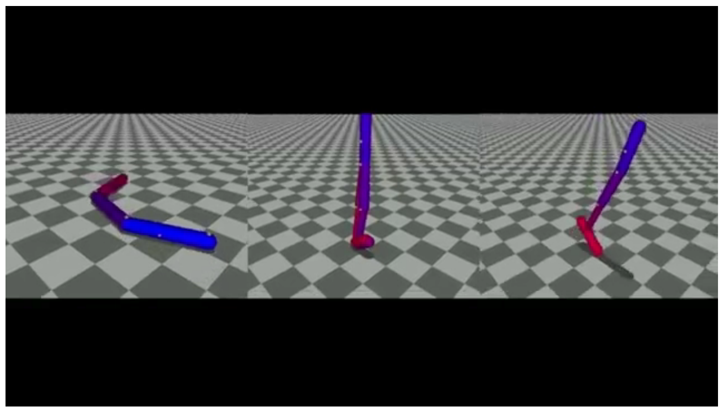

[Schulman et al. 2016]

[Chebotar et al. 2017]

REINFORCEMENT LEARNING FOR LOCOMOTION CONTROL

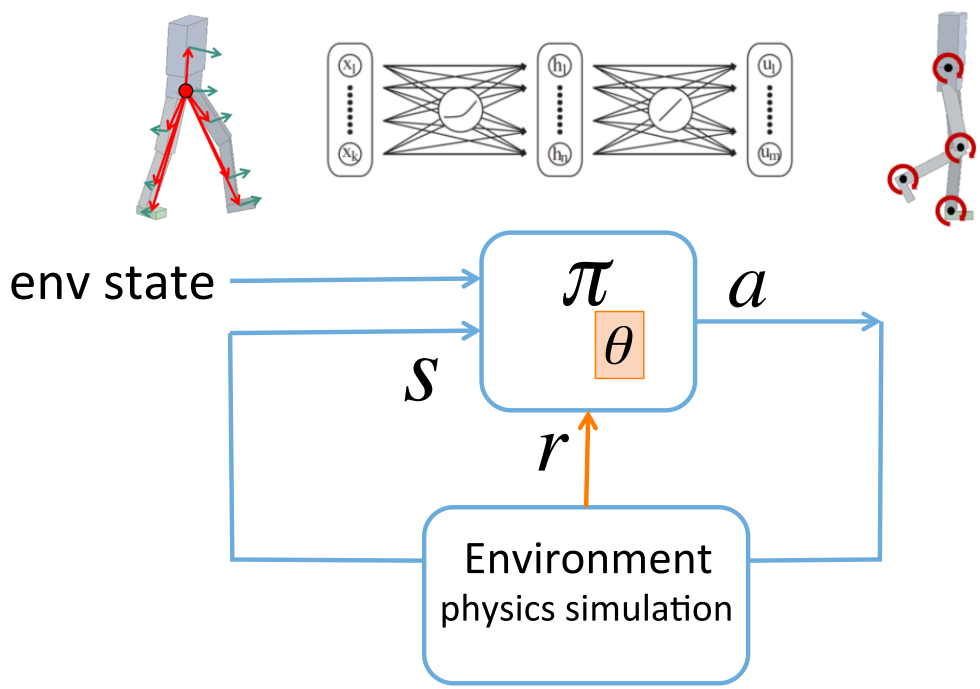

In principle:
- specify rewards
- “train” using RL algorithm

$$\max_{\theta} \quad \mathrm{E}[\sum_{t=0}^{H} R(s_t)|\pi_\theta]$$

MOTION QUALITY

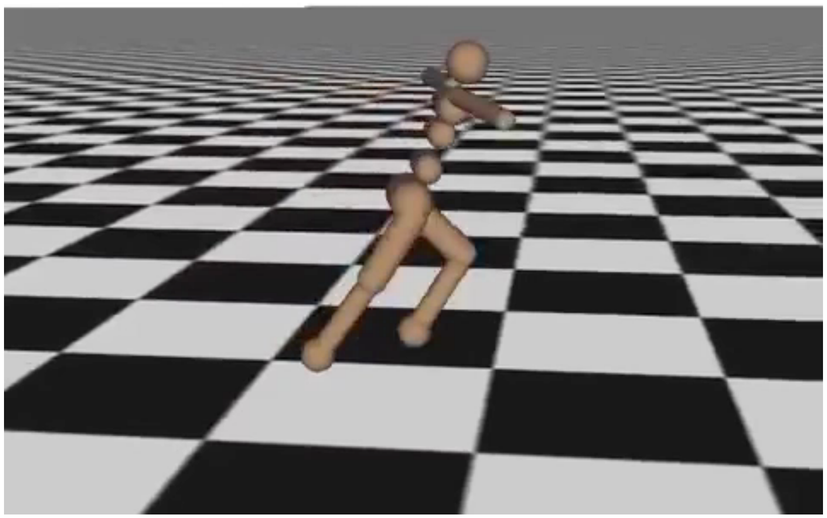

[Schulman et al. 2016]

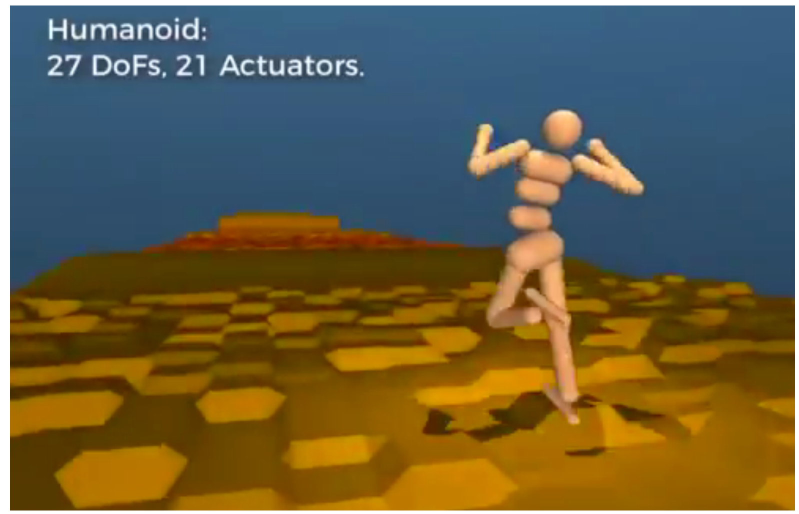

[Heess et al. 2017]

IN PRACTICE

network structure

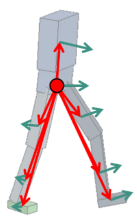

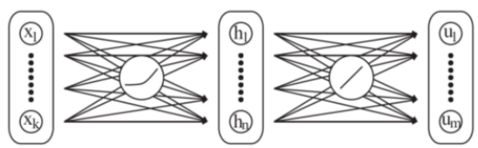

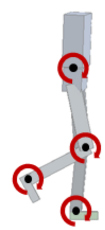

structure,
torque limits,
friction,
…

initial states
noise amplitude
batch size
step-size control
early termination
learning iterations
…

state
description

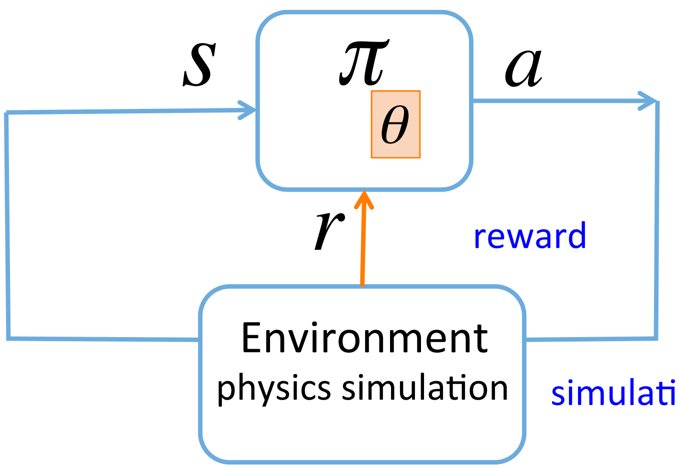

control dt

reward

simulation dt

- how to specify rewards?
- inefficient?
- local minima?
- poor motion quality?
- sim2real?
- reproducible?

OVERVIEW

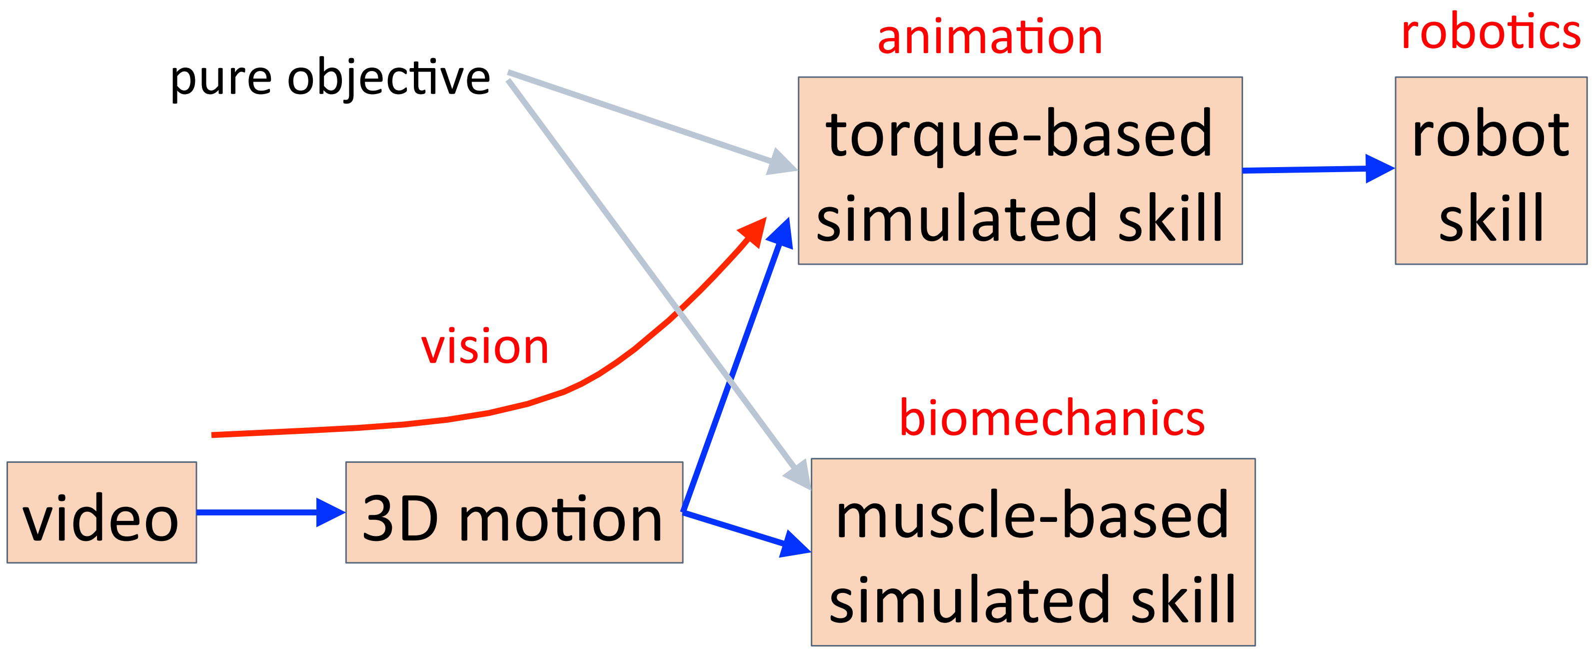
pure objective
animation
robotics
torque-based simulated skill
robot skill
vision
biomechanics
video
3D motion
muscle-based simulated skill

MOTION IMITATION

$$r_t = \omega^I r_t^I + \omega^G r_t^G$$

Imitation Objective

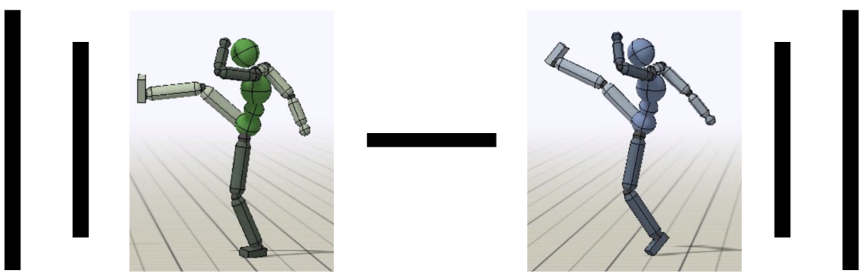

Task Objective

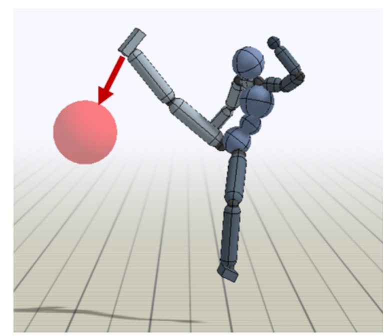

STATE + ACTION

State: 197 D

- link positions
- link velocities

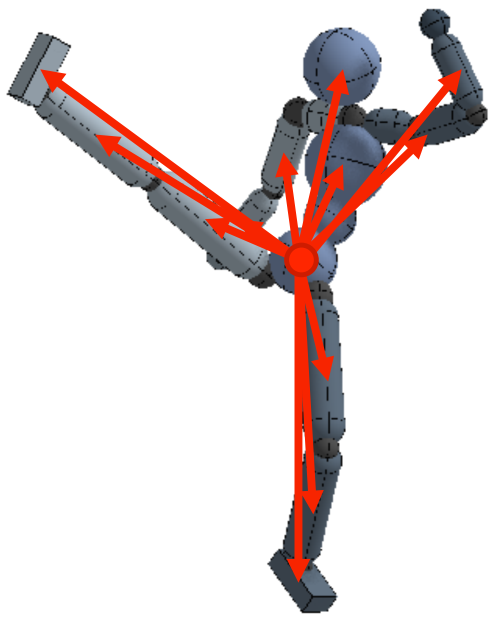

Action: 36 D

- PD targets

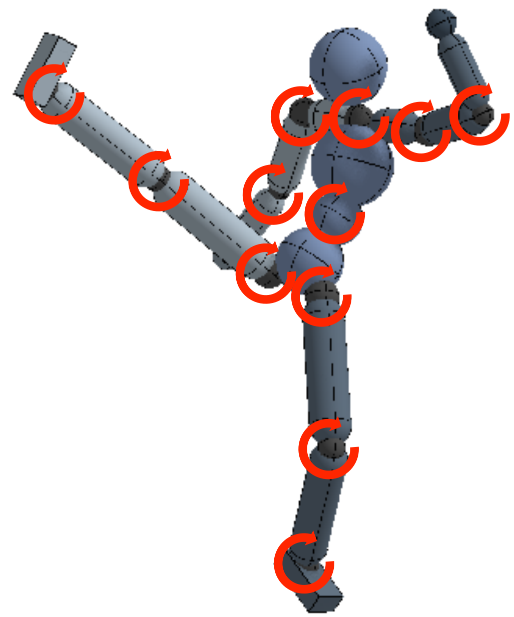

WALKING

[DeepLoco: SIGGRAPH 2017]

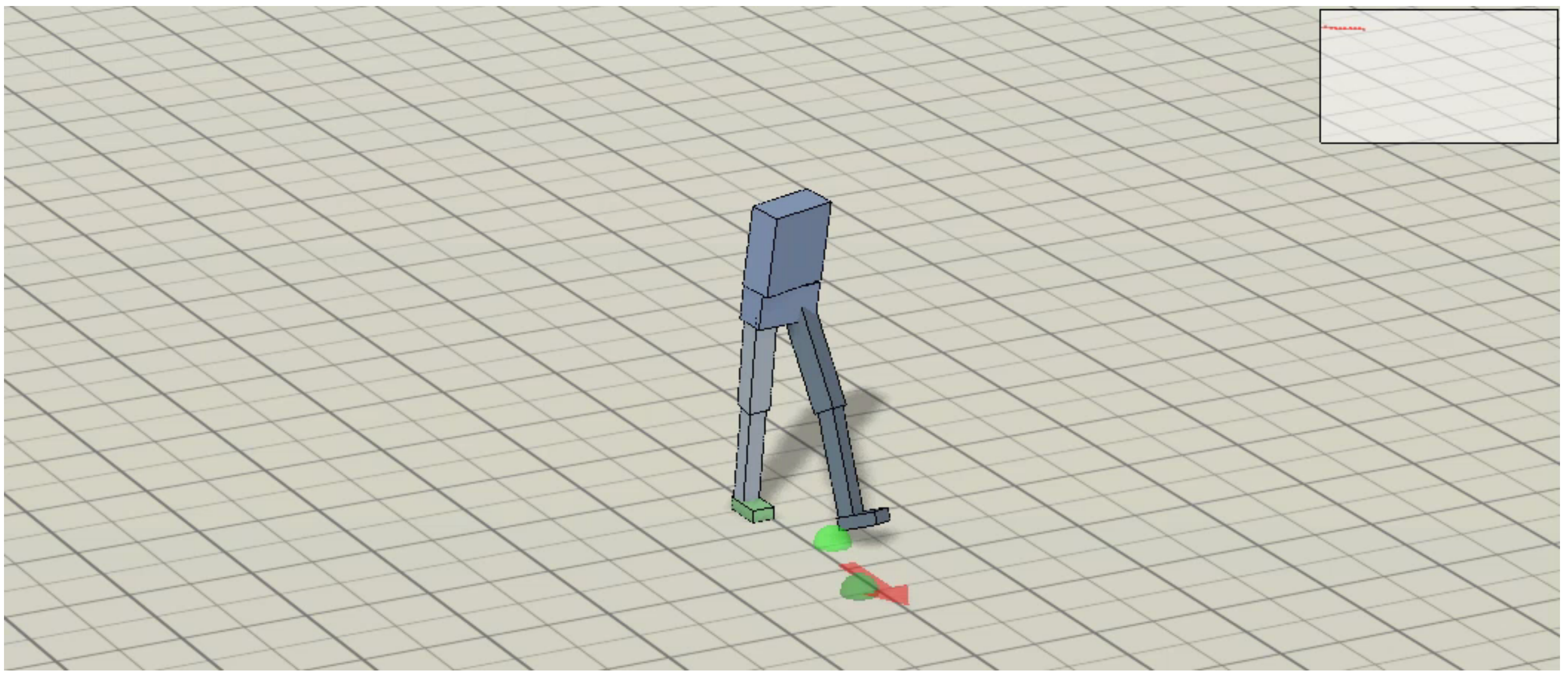

[DeepLoco: SIGGRAPH 2017]

Walking on Conveyor Belts

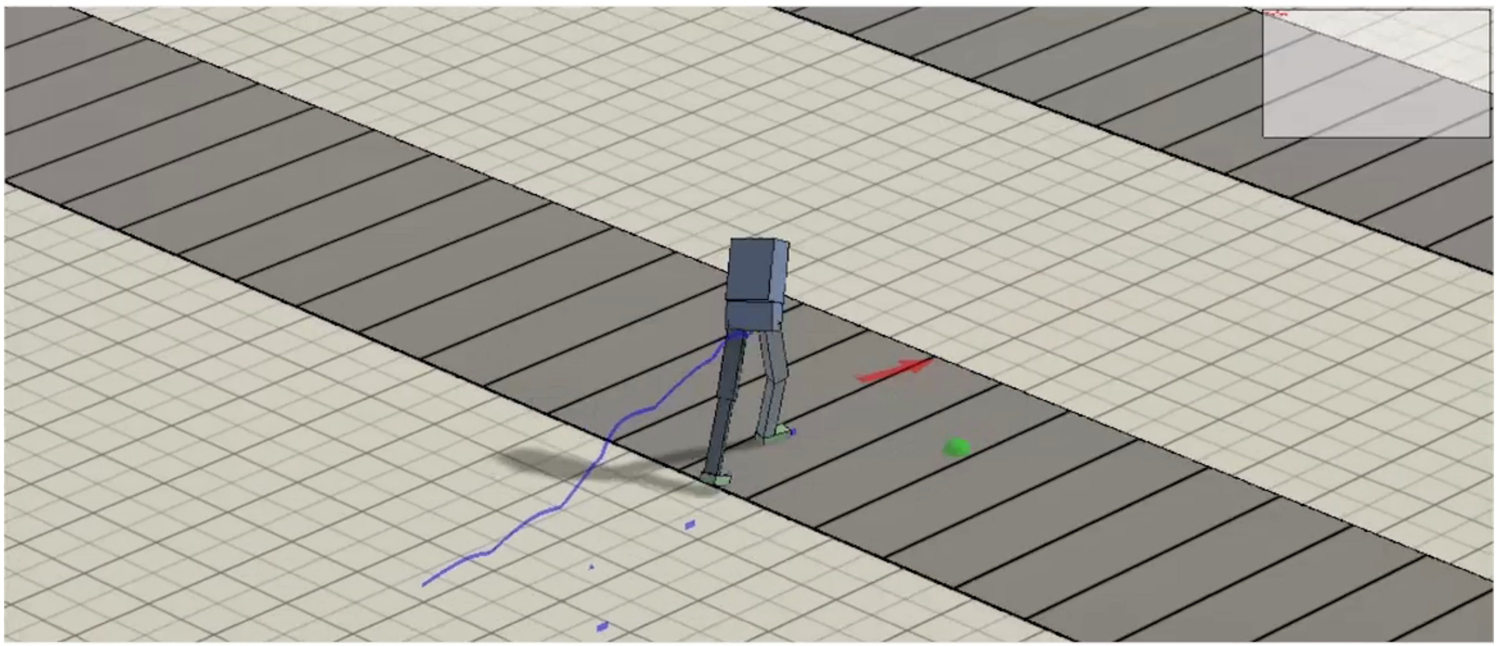

Forward Lean

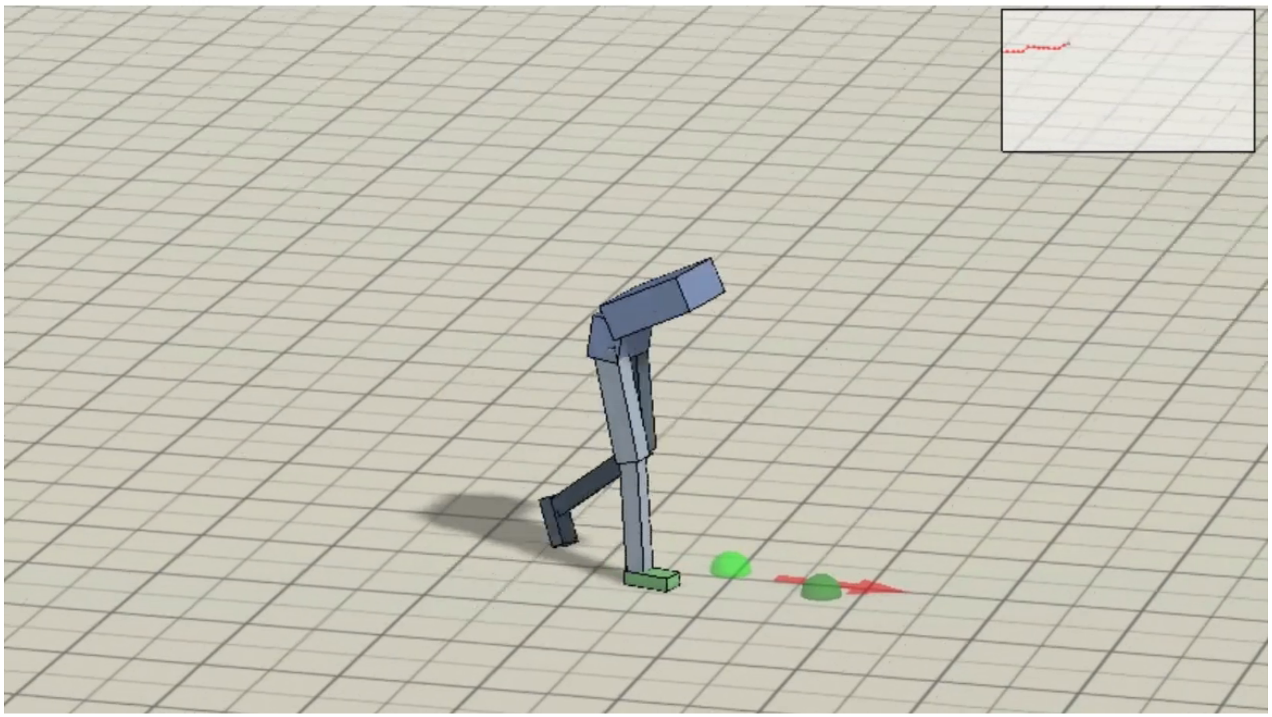

STATE + ACTION

State: 197 D

- link positions
- link velocities

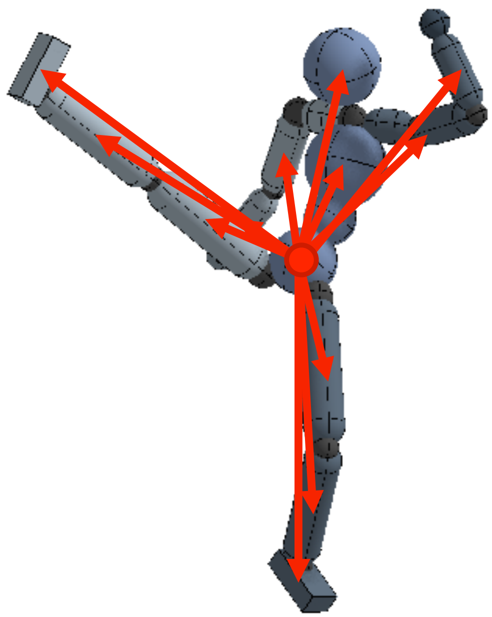

Action: 36 D

- PD targets

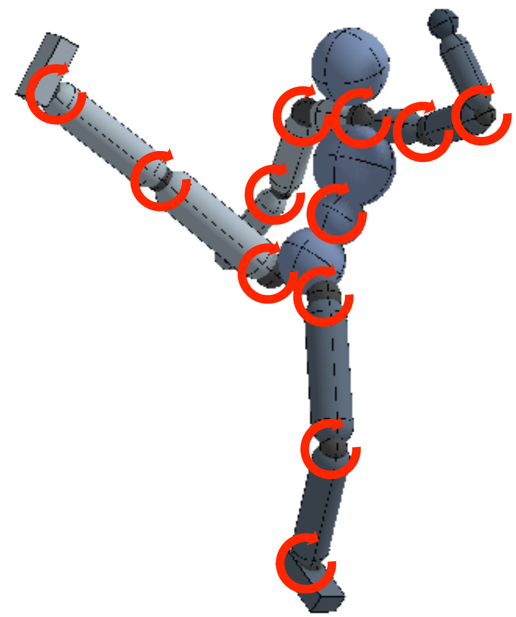

NO REFERENCE MOTION

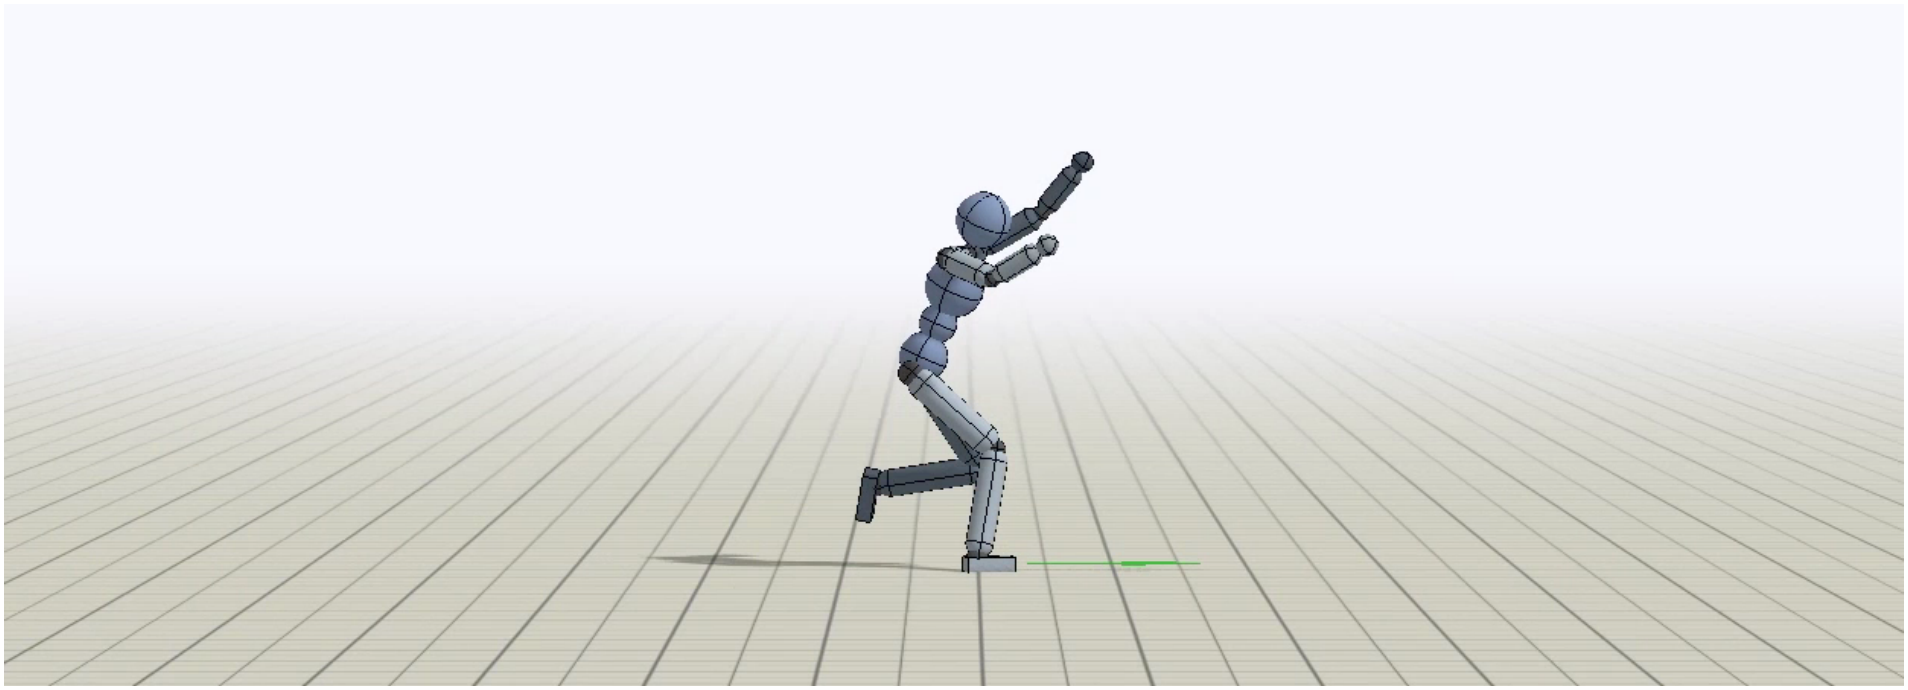

HUMANOID: RUN

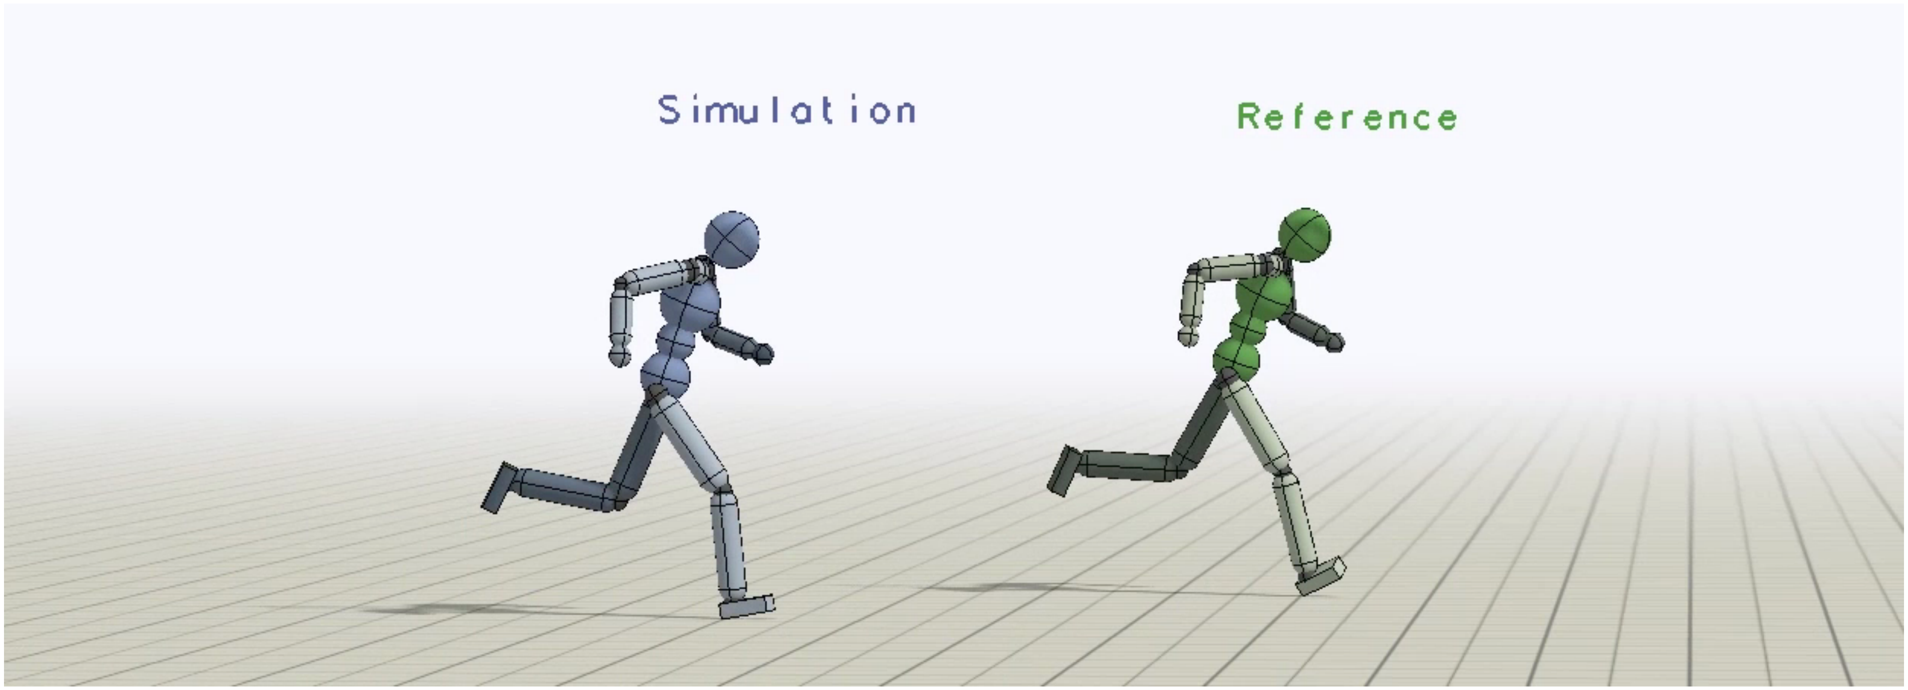

HUMANOID: BACKFLIP

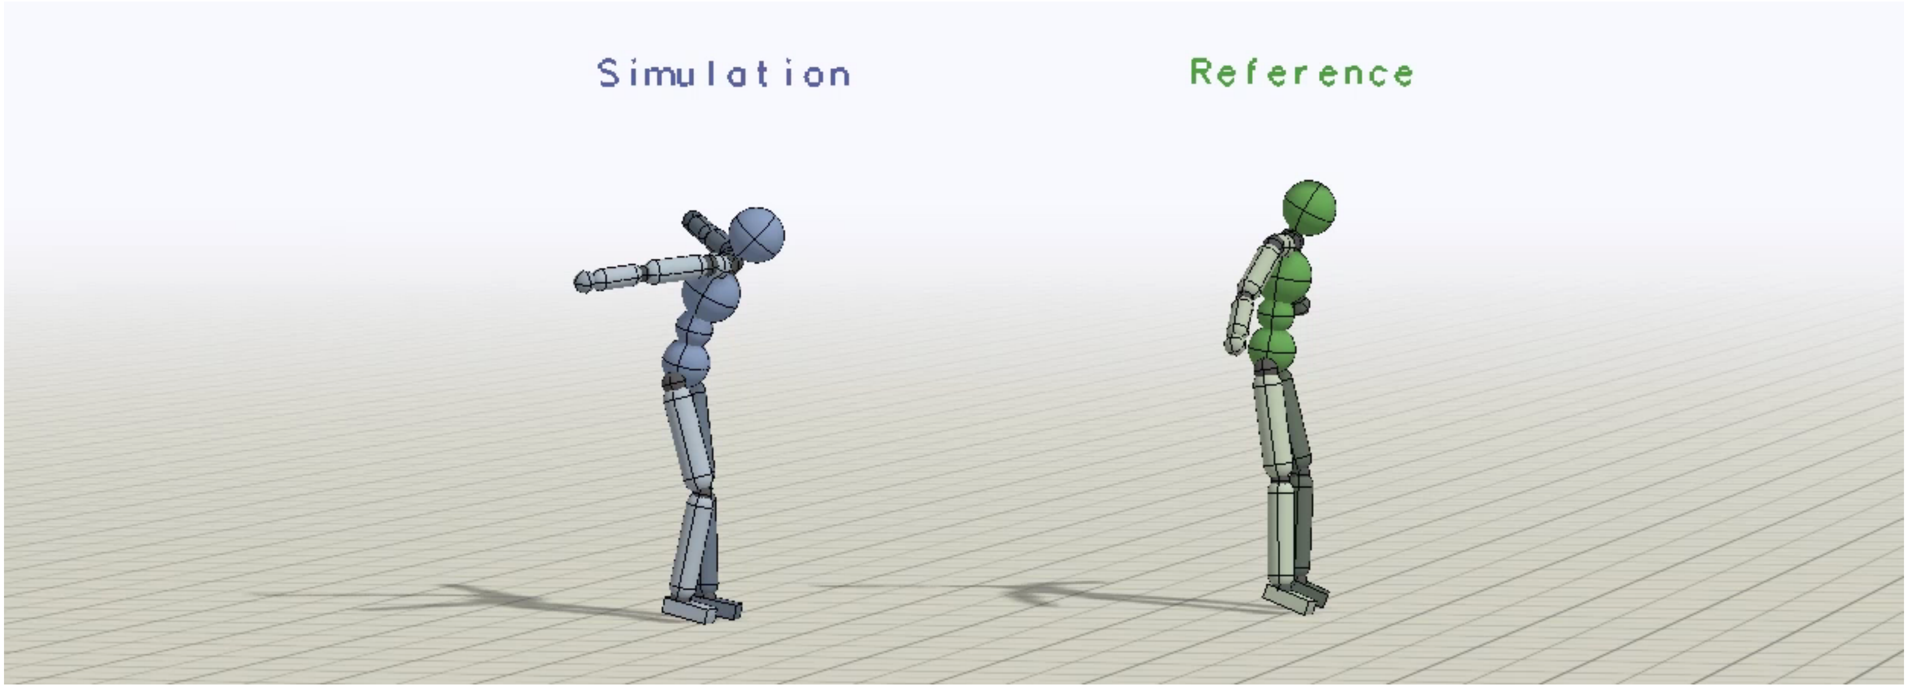

LOCAL MINIMA

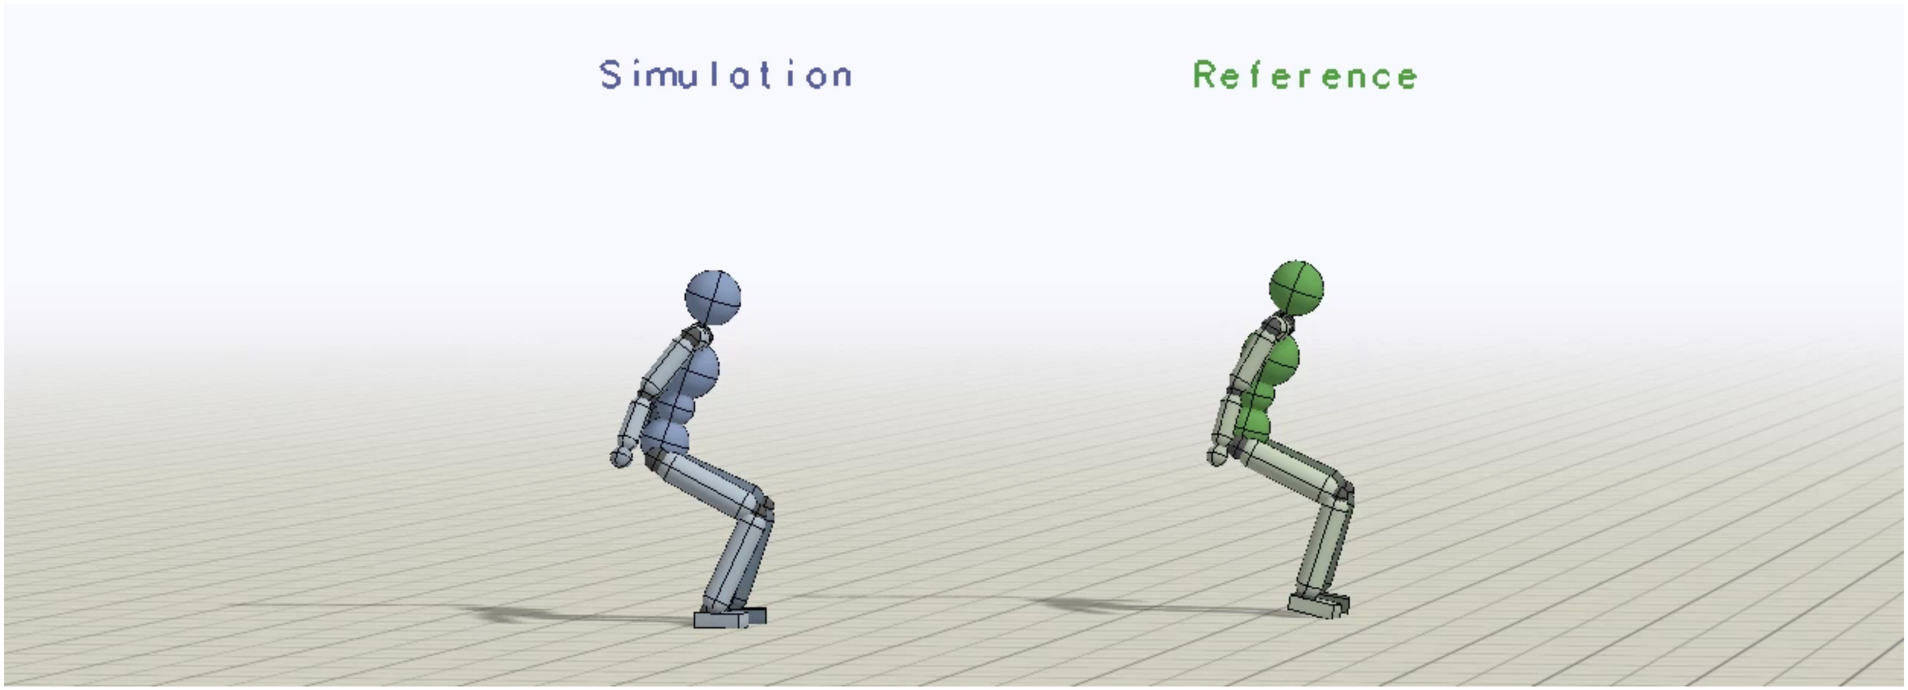

WITH REFERENCE STATE INITIALIZATION

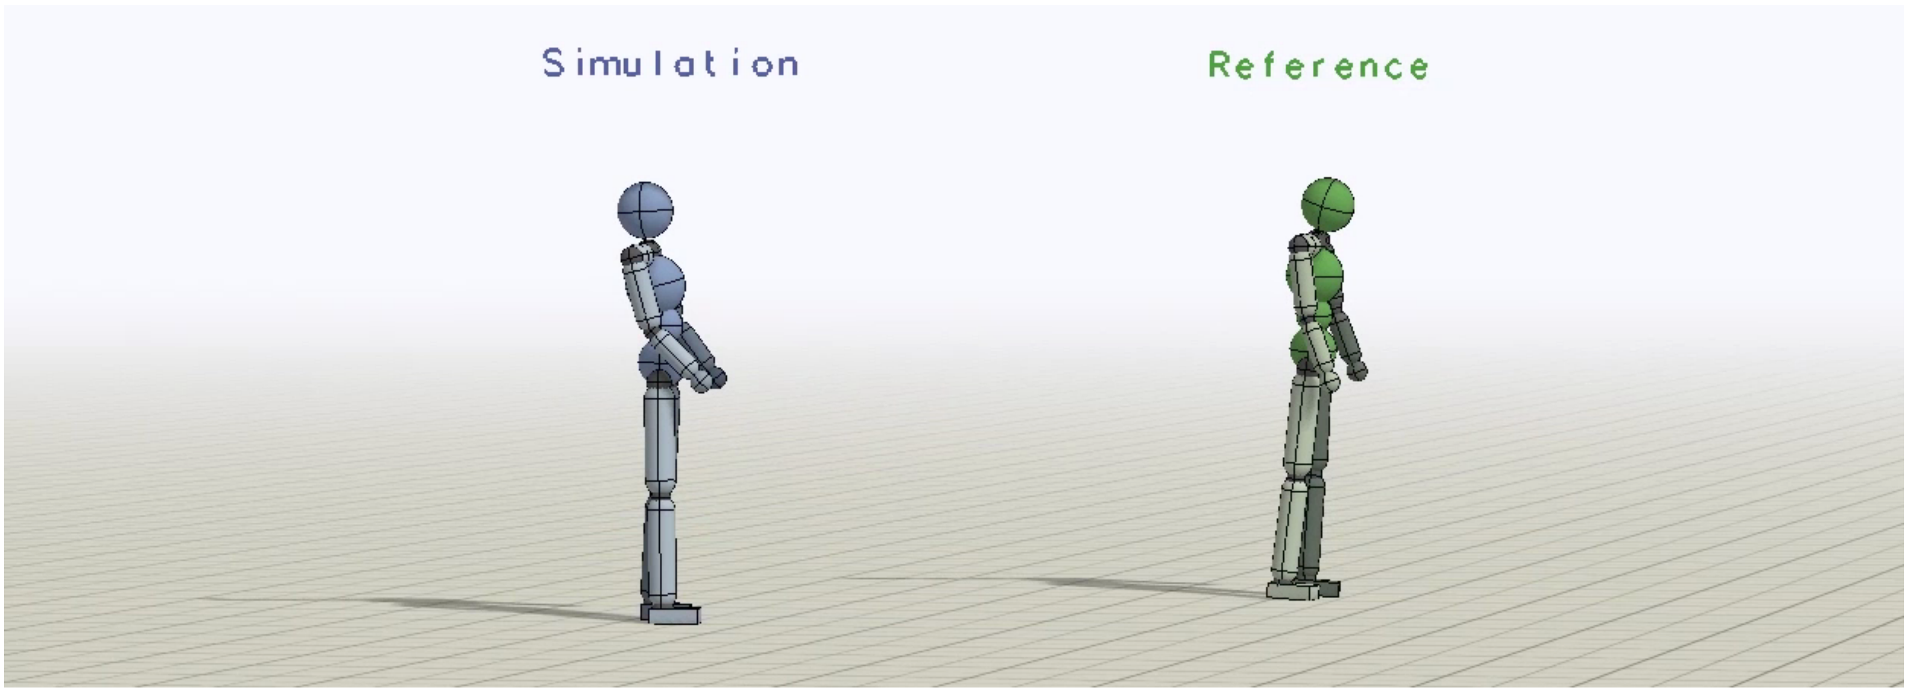

MULI-CLIP INTEGRATION

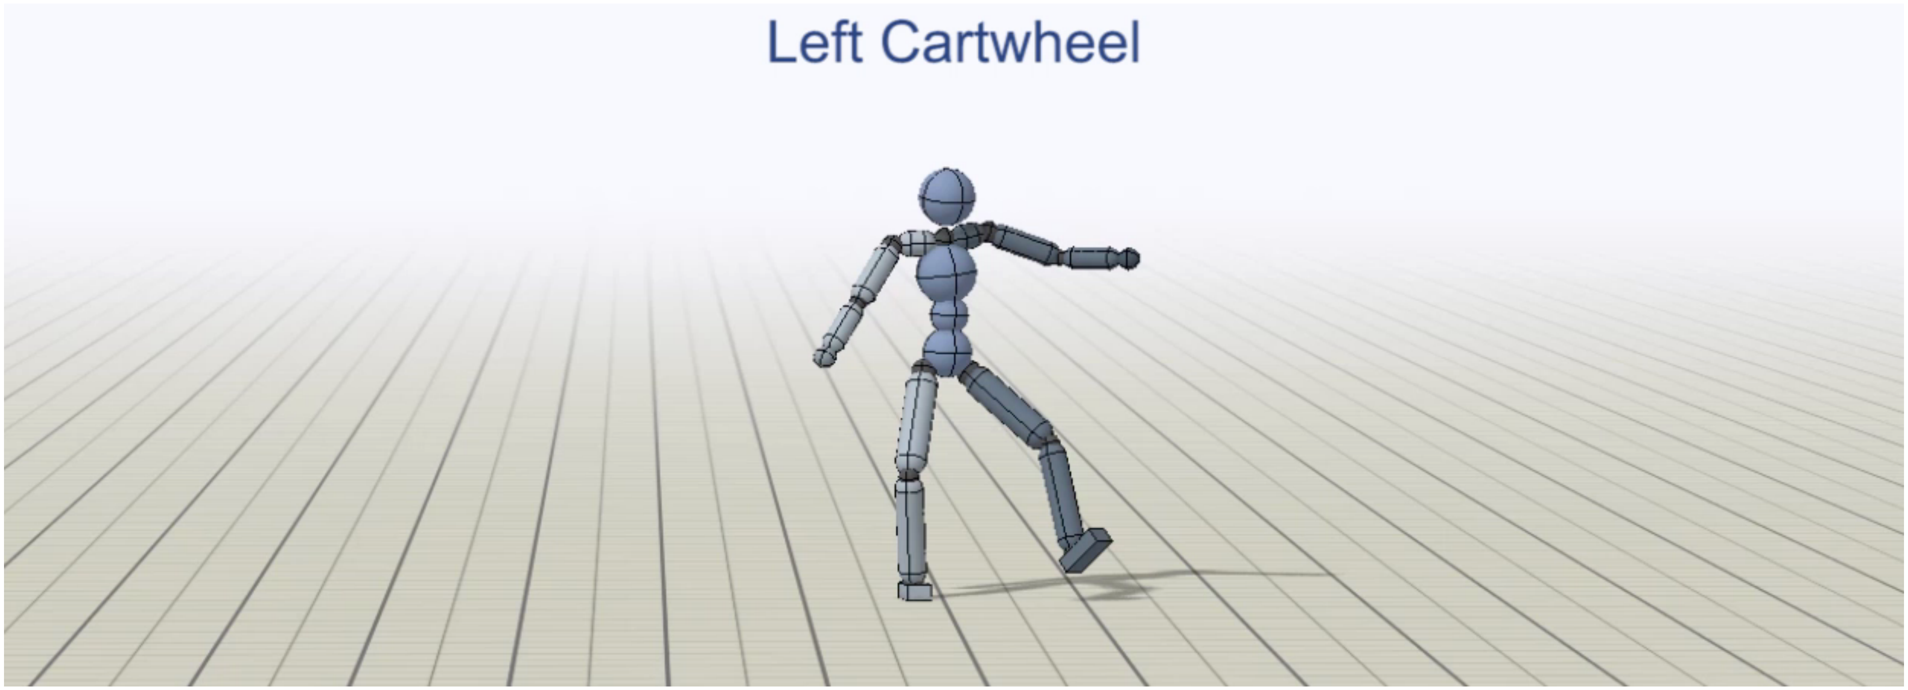

HUMANOID: BALANCE BEAM

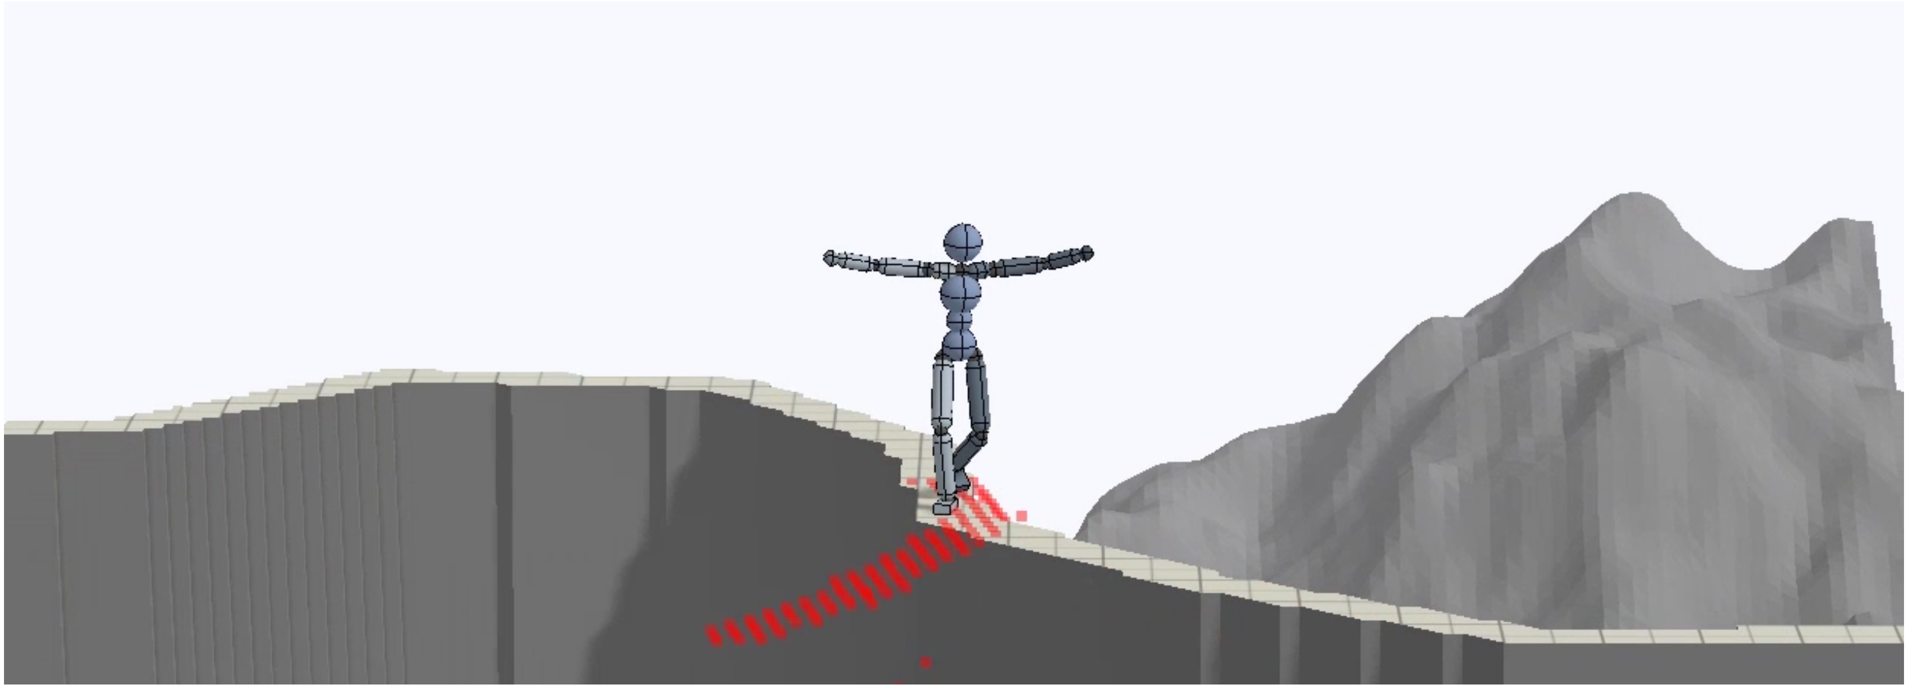

HUMANOID: RUN – DENSE GAPS

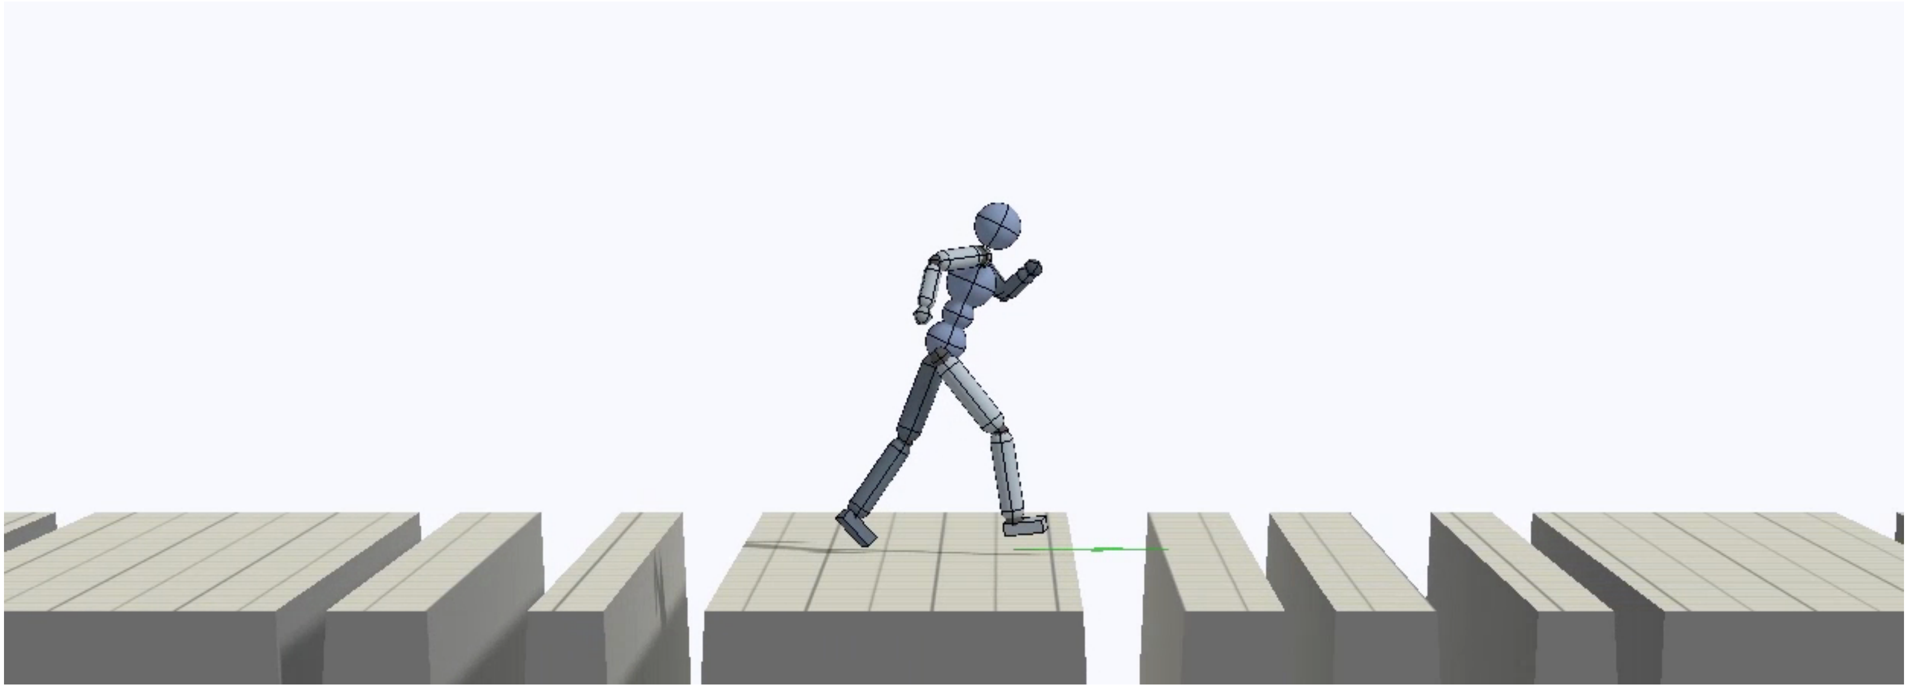

CHARACTER RETARGETING

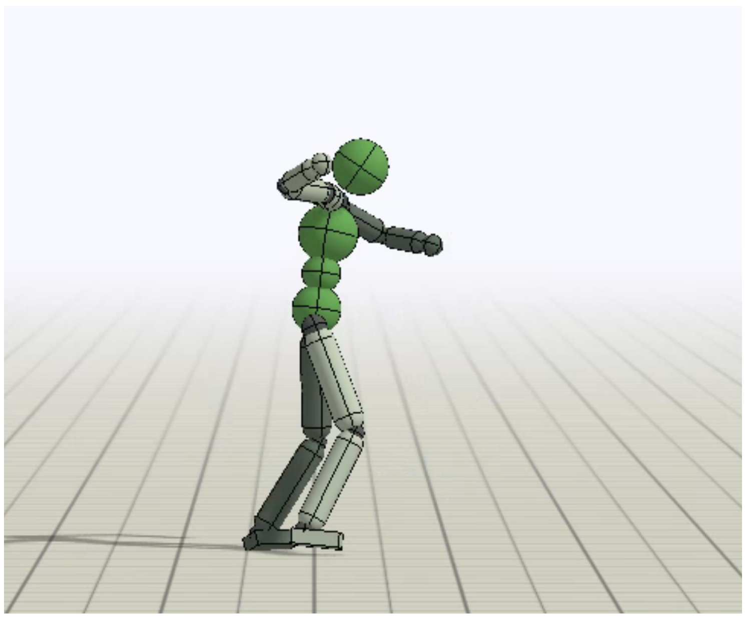

Reference Motion

+

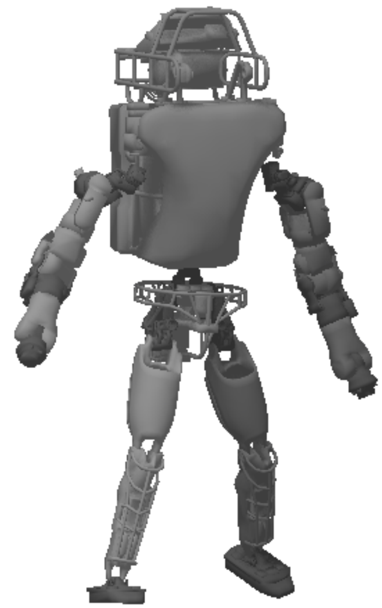

Atlas

ATLAS: SPINKICK

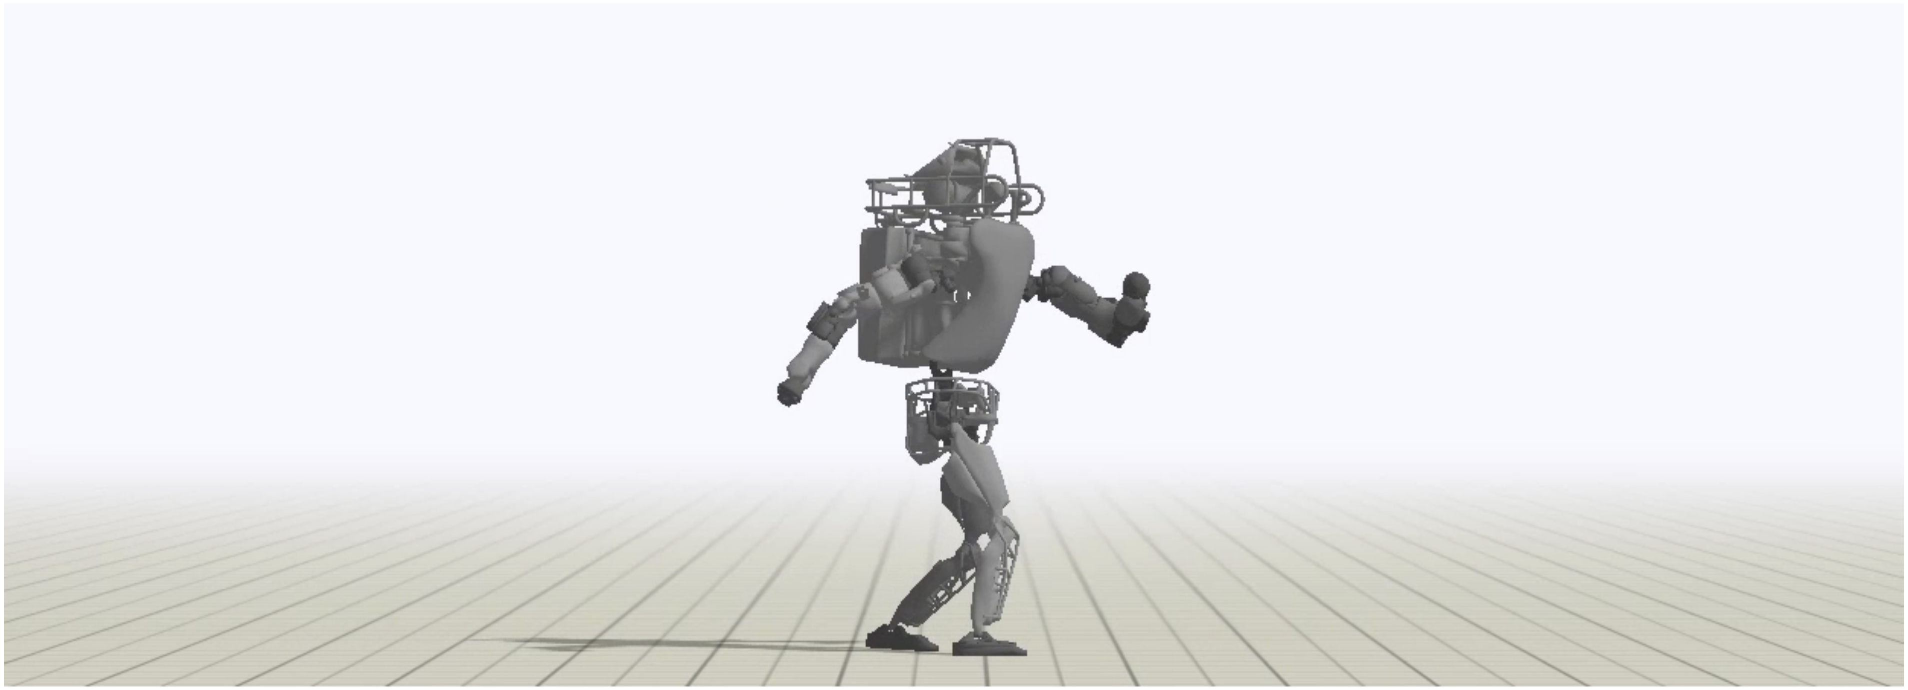

ATLAS: GETUP-FACEDOWN

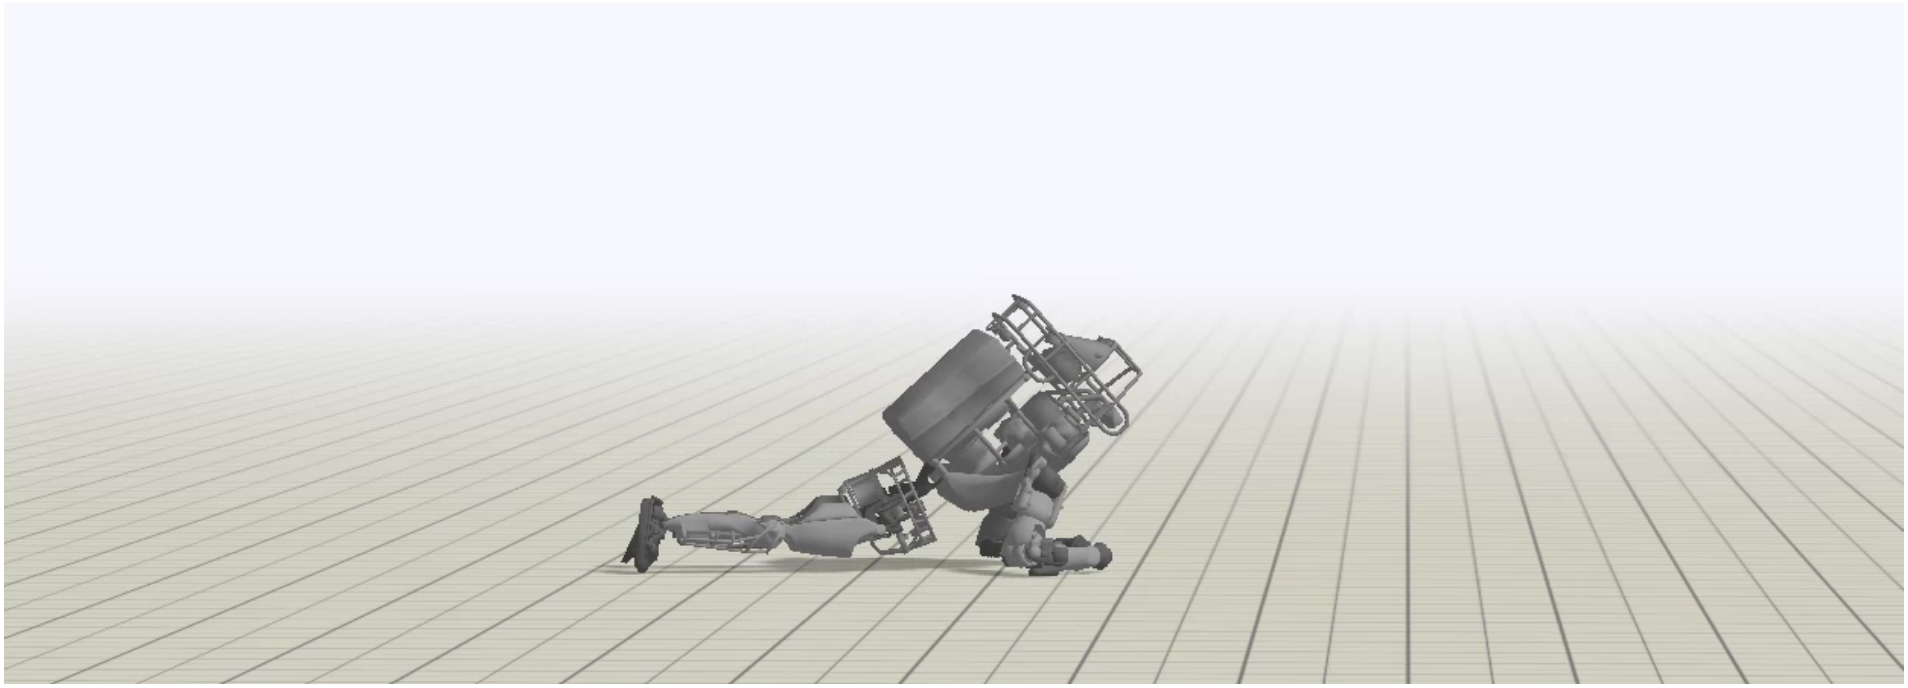

SIMULATED LION

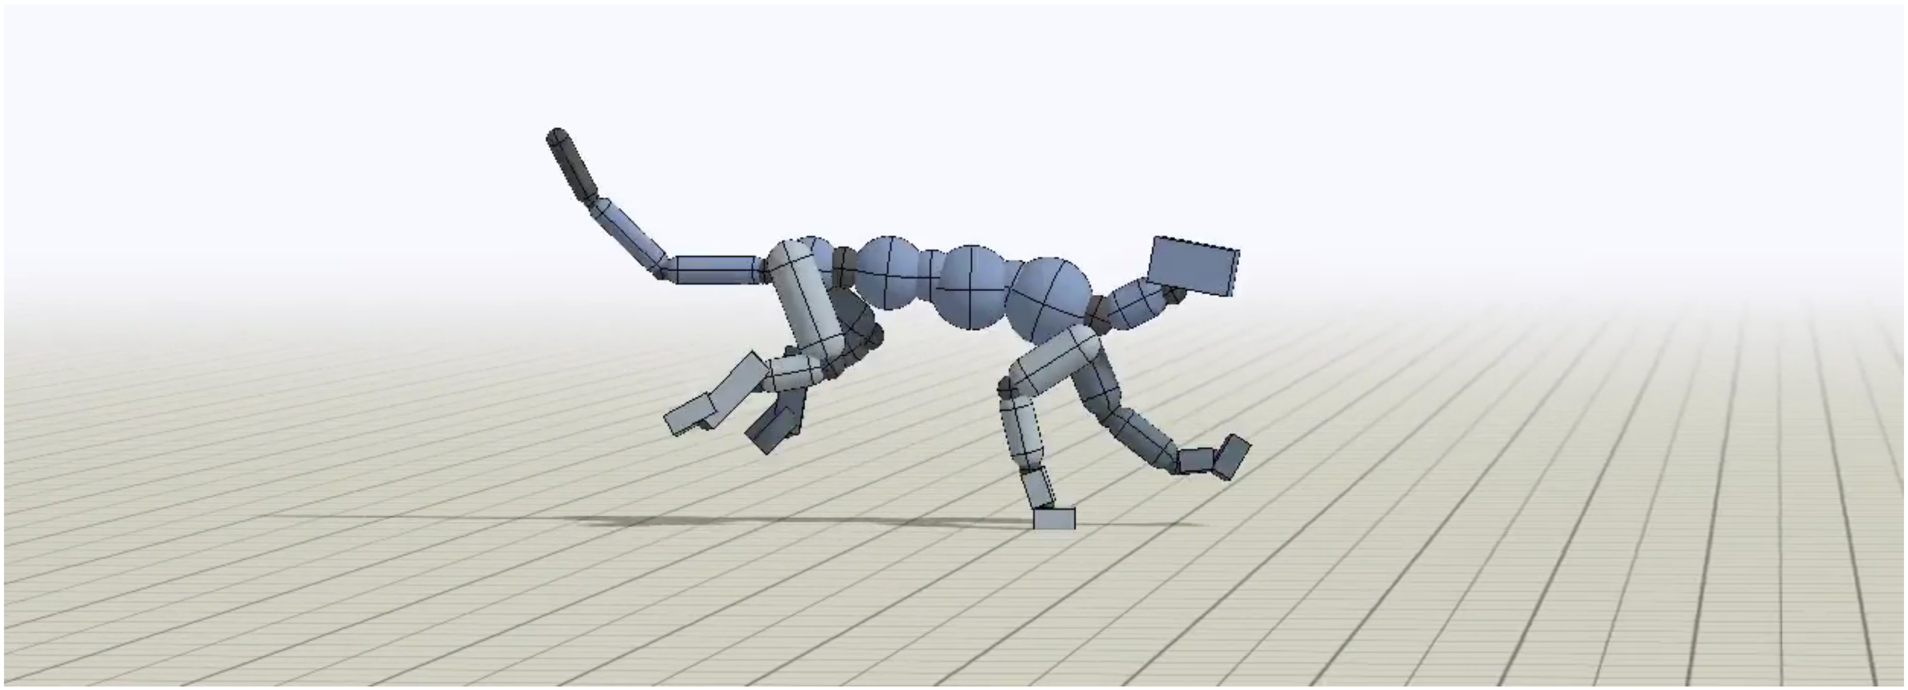

Ziva Dynamics

ZIVA


ZIVA

DEEPLOCO: HIERARCHICAL RL

[SIGGRAPH 2017]

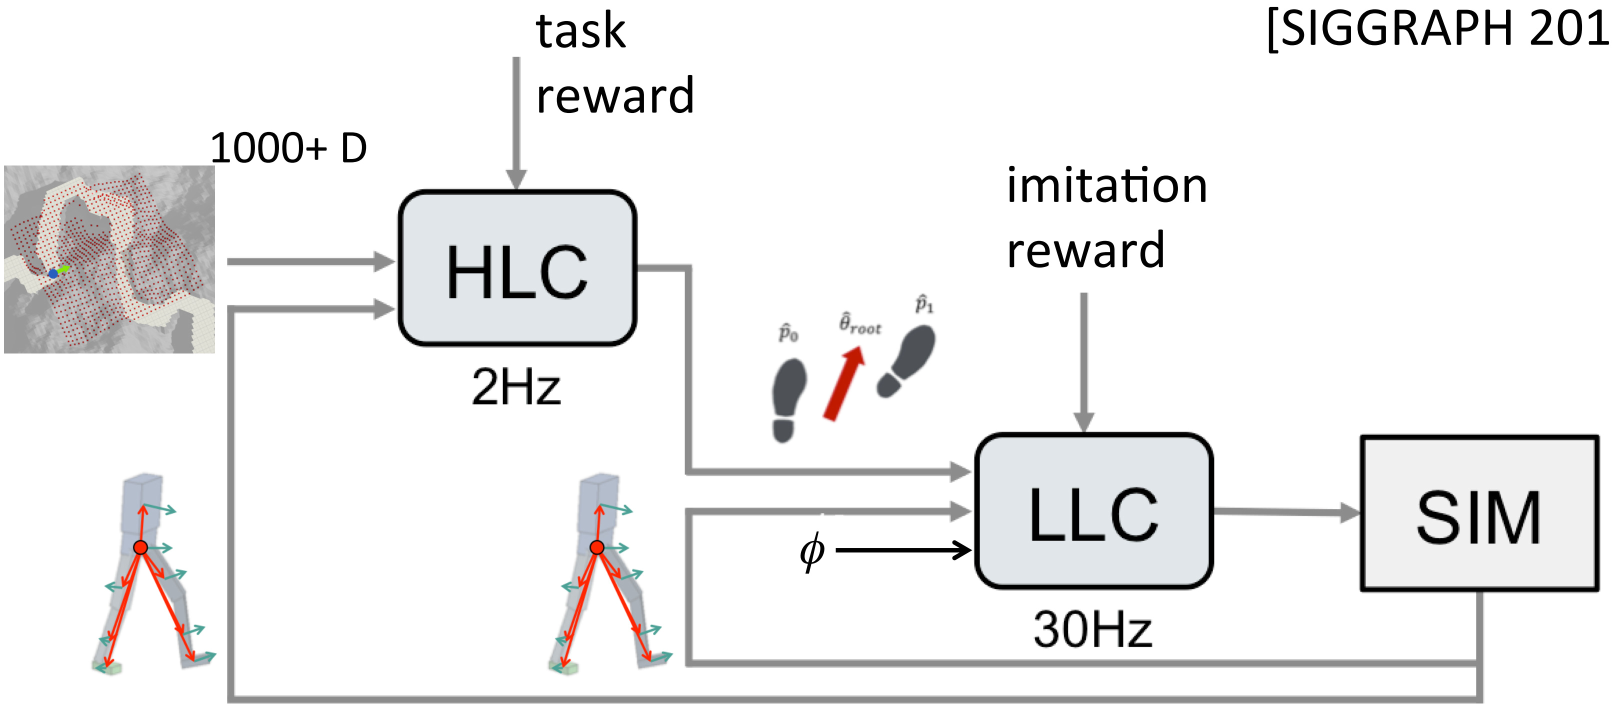

Dynamic Obstacles

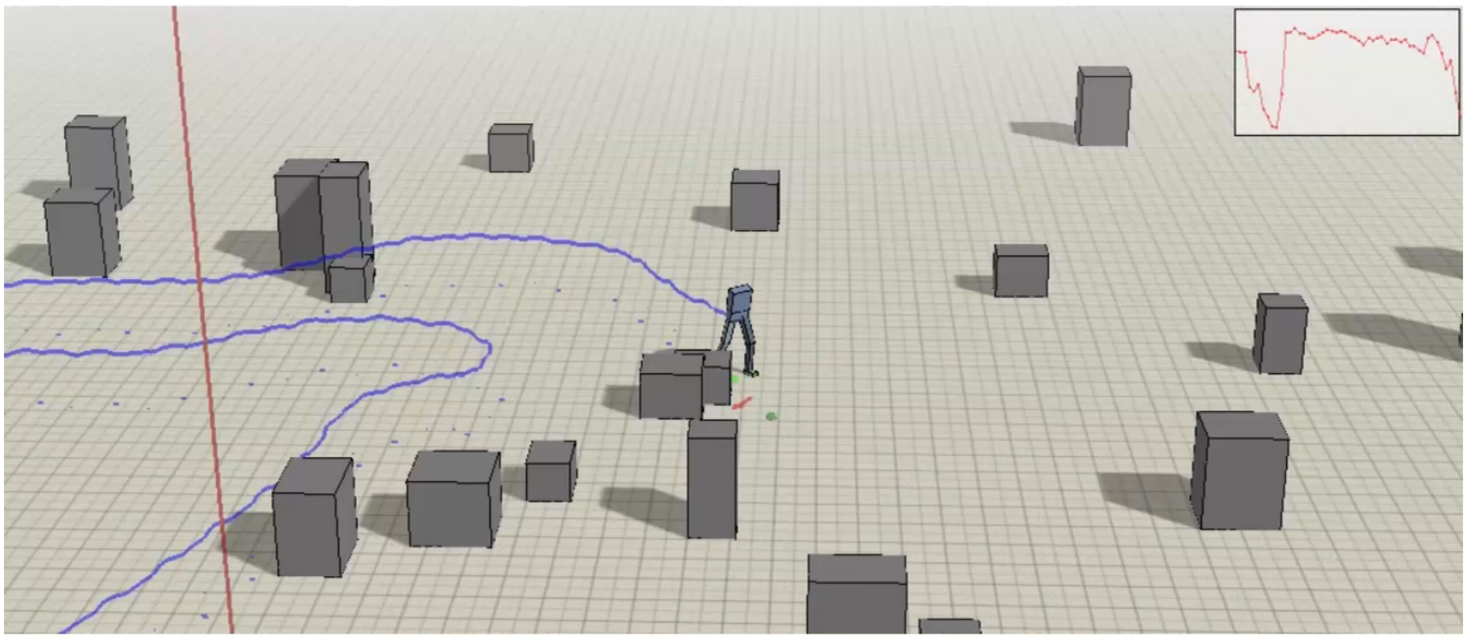

HLC

Soccer Dribbling

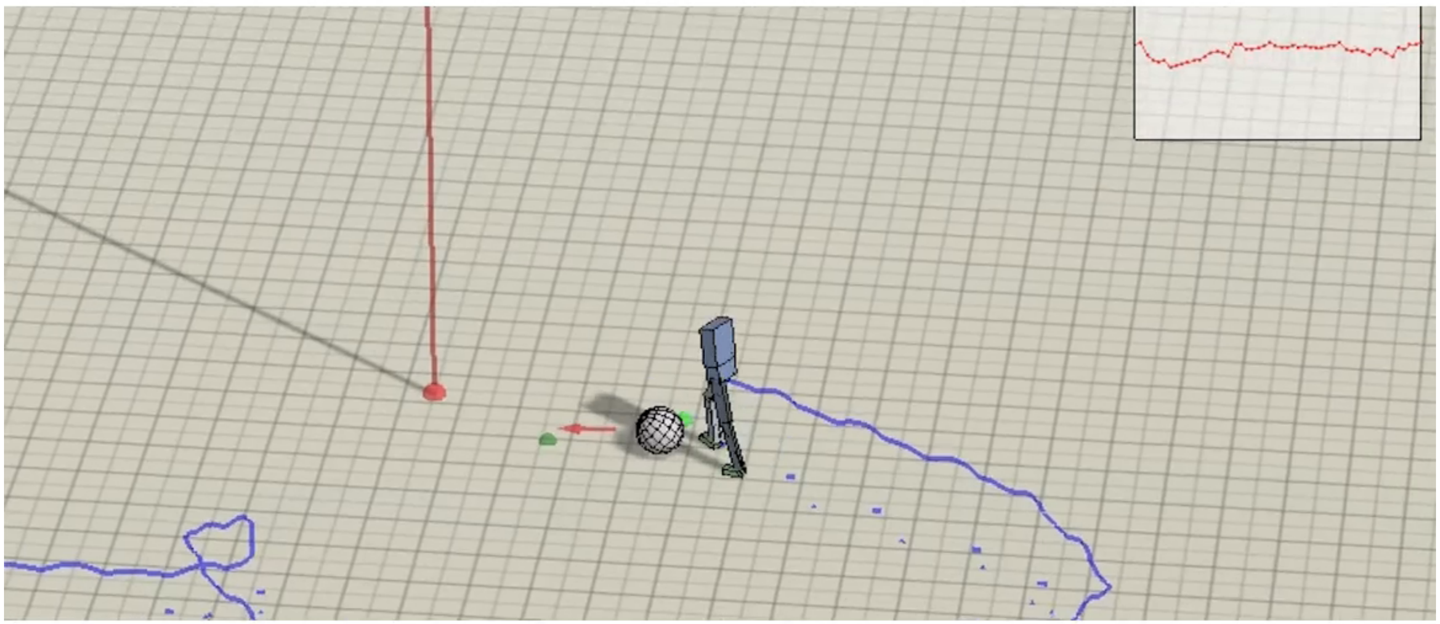

SKILLS FROM VIDEO:
REINFORCEMENT LEARNING OF PHYSICAL SKILLS FROM VIDEO

[SIGGRAPH ASIA 2018]

Transactions on Graphics (Proc. ACM SIGGRAPH Asia 2018)

Xue Bin Peng Angjoo Kanazawa Jitendra Malik Pieter Abbeel Sergey Levine
University of California, Berkeley

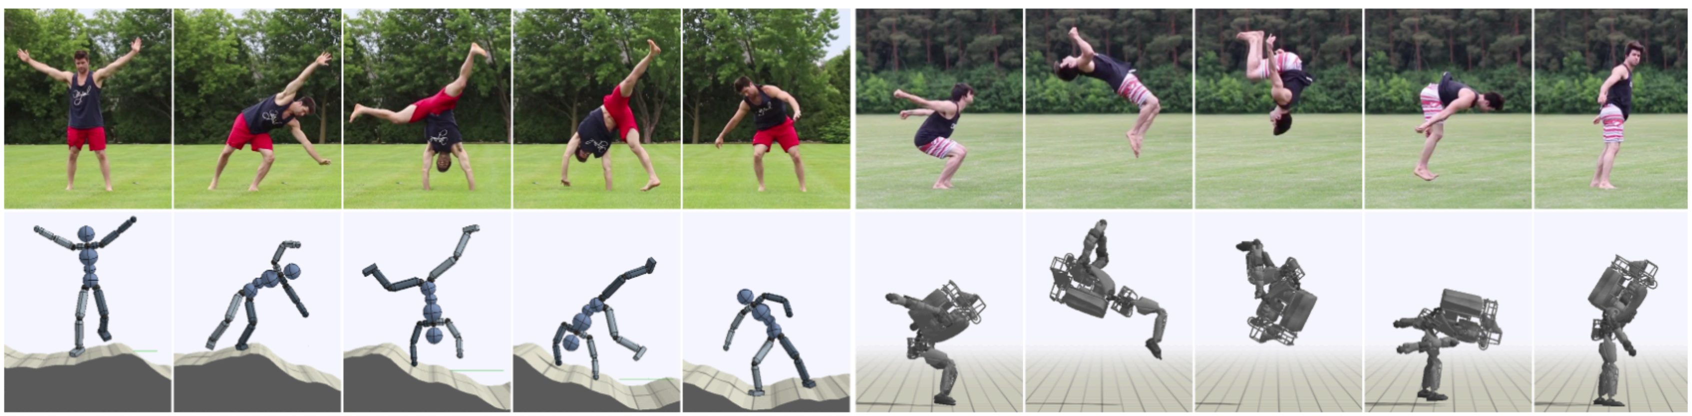

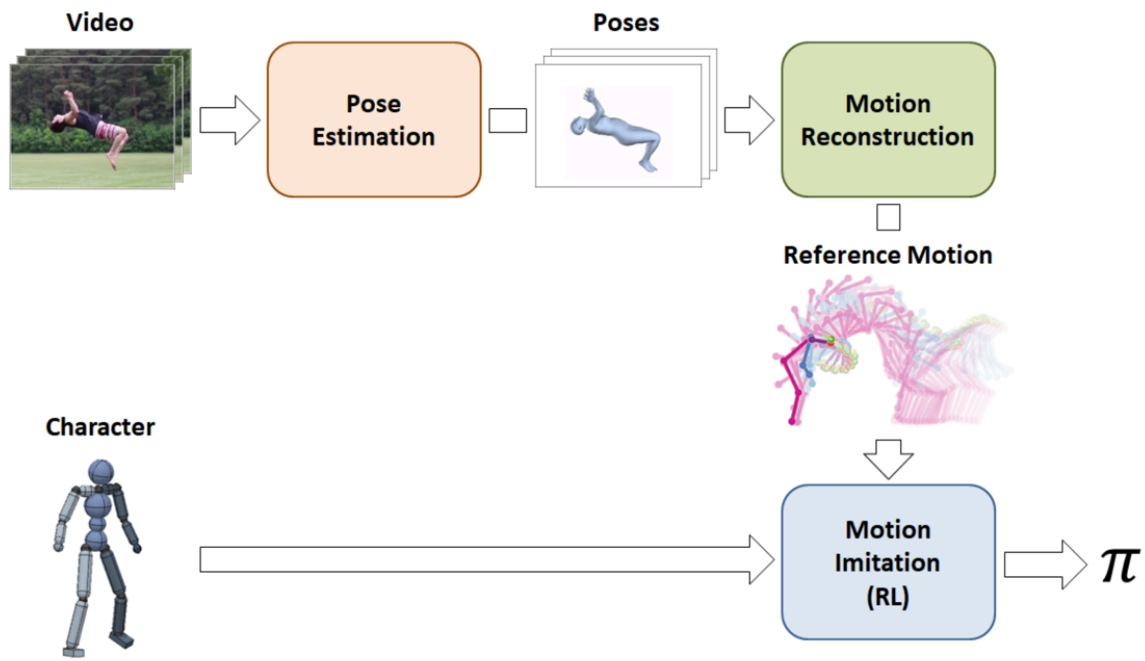
Video
Pose
Estimation
Poses
Motion
Reconstruction
Reference Motion
Character
Motion
Imitation
(RL)
π

SKILLS FROM VIDEOS

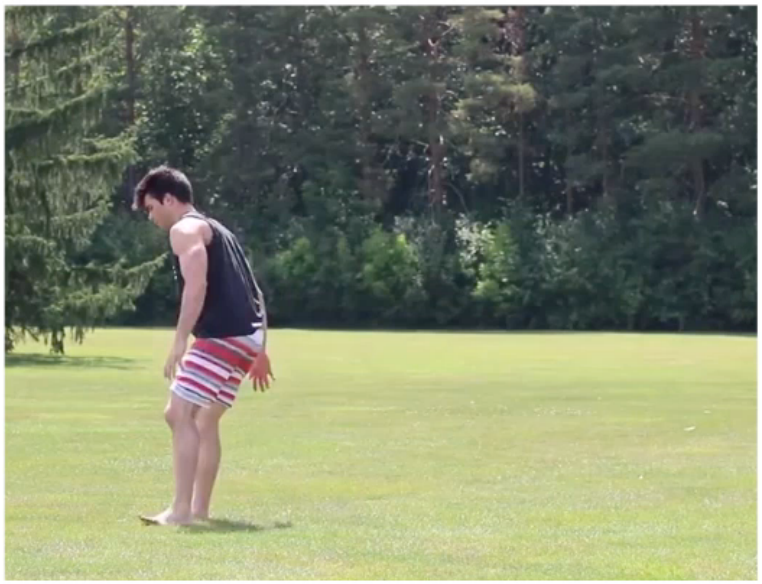

Video

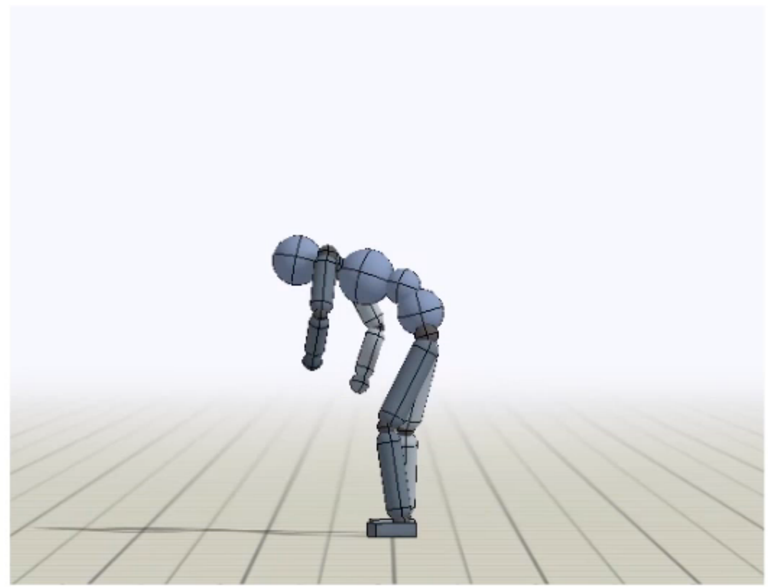

Simulation

SKILLS FROM VIDEOS

Video

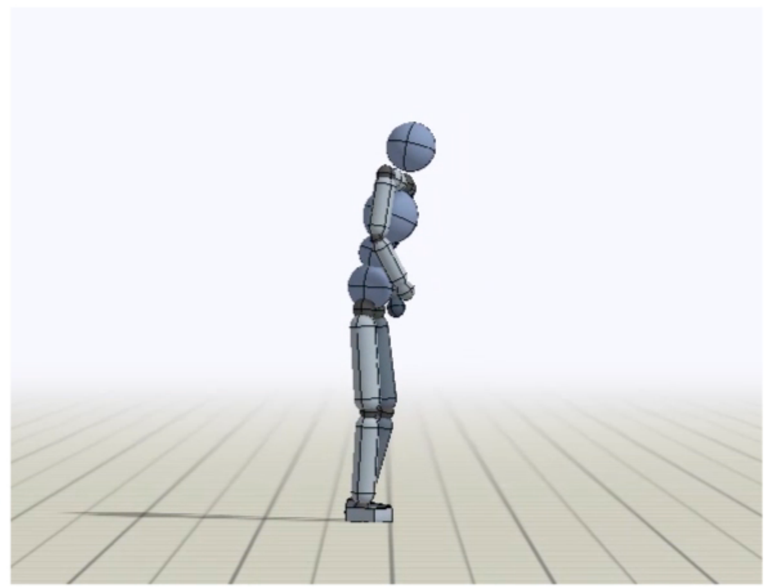

Simulation

Retargeting

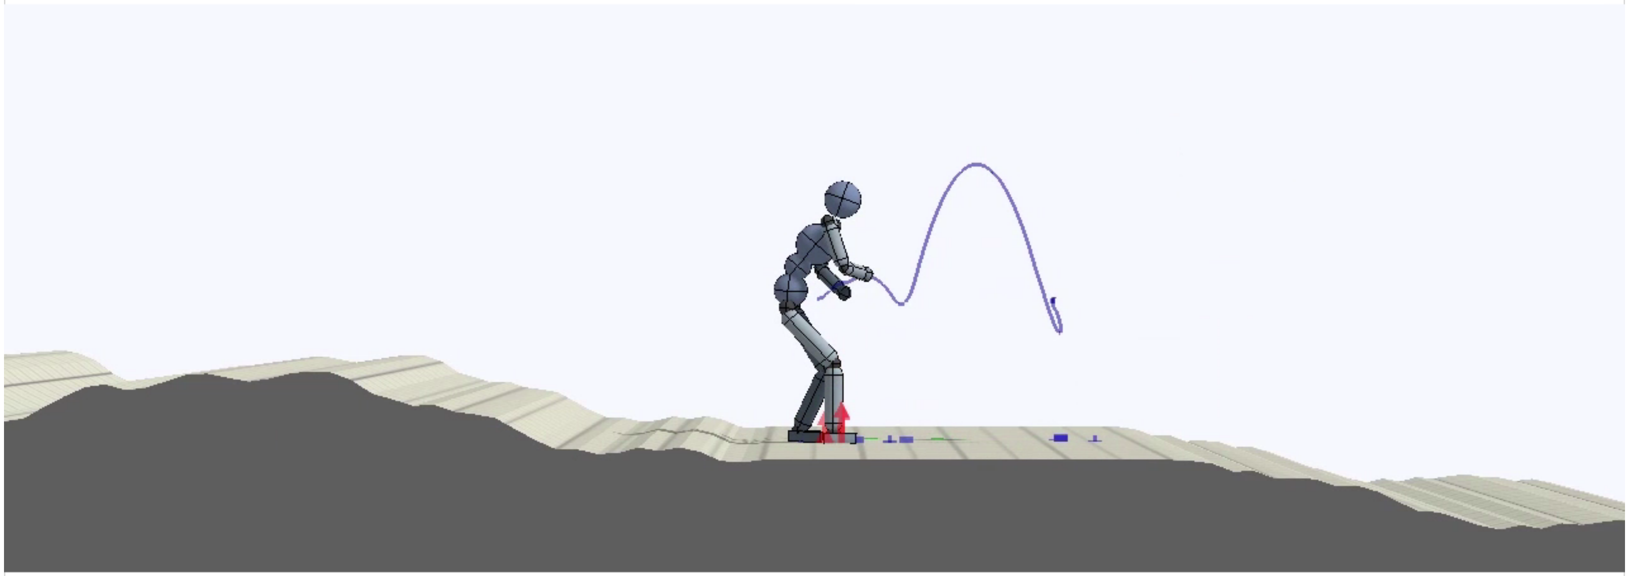

and retarget to different environments.

Character Retargeting

Video: Jump

Policy

MOTION COMPLETION

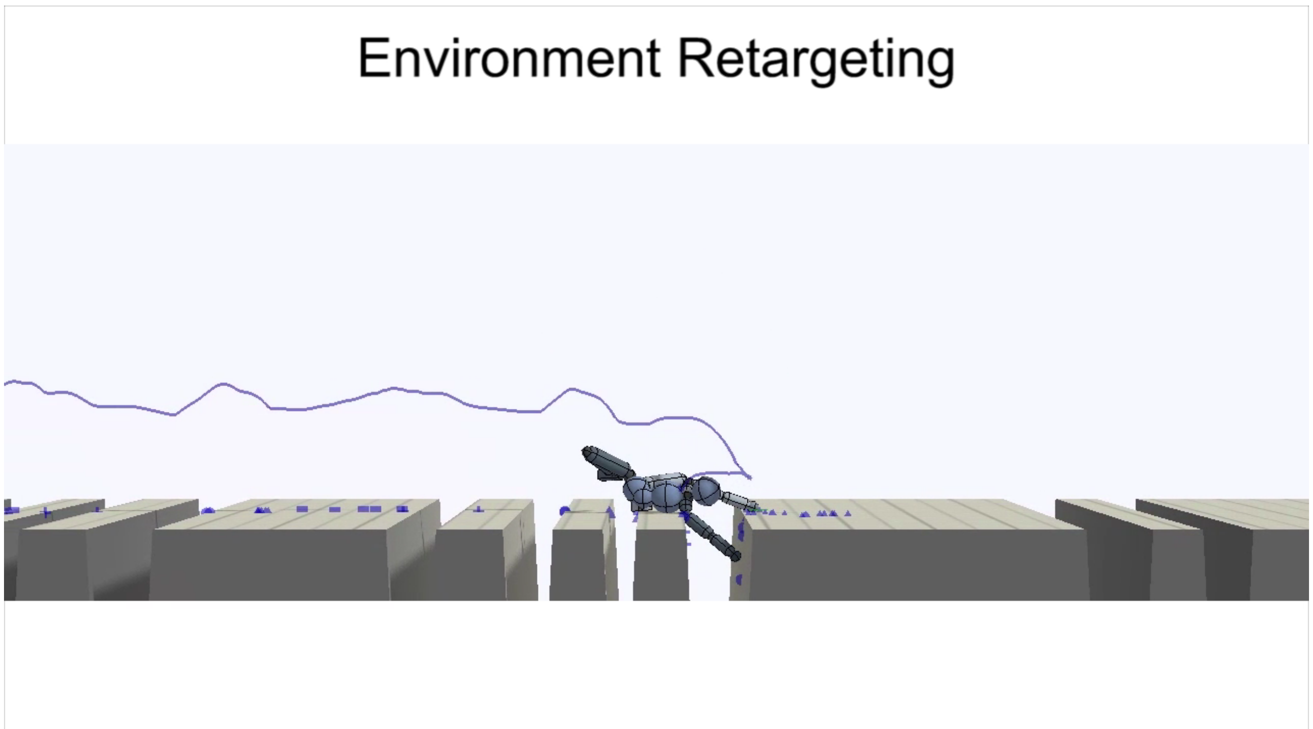

FEEDBACK CONTROL FOR CASSIE WITH DEEP REINFORCEMENT LEARNING

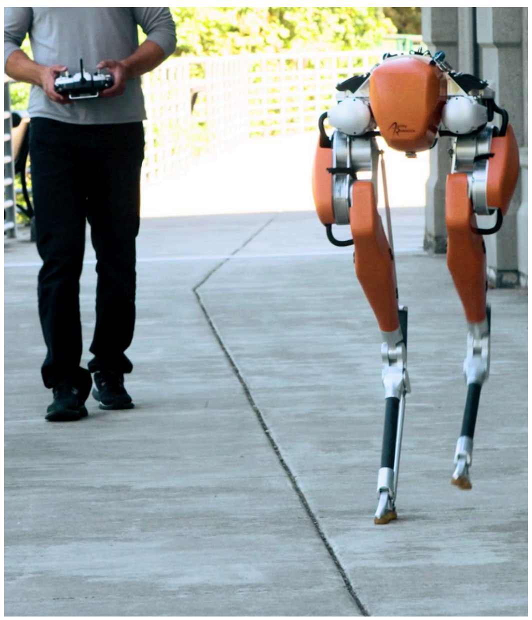

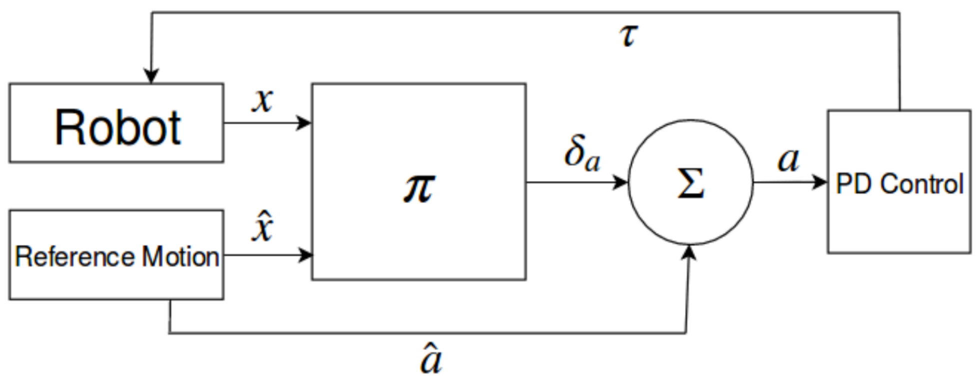

WALK AND STUMBLE

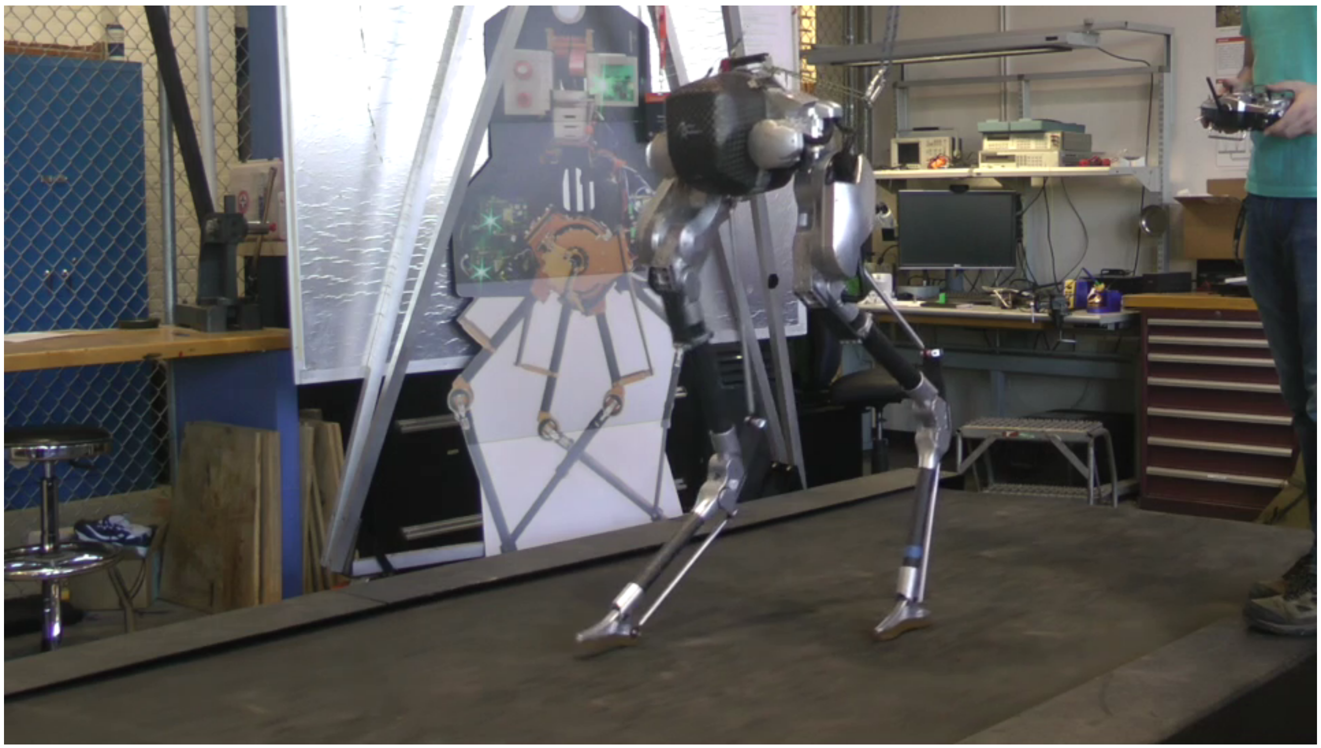

BACKWARDS WALK

SIDE-STEPS

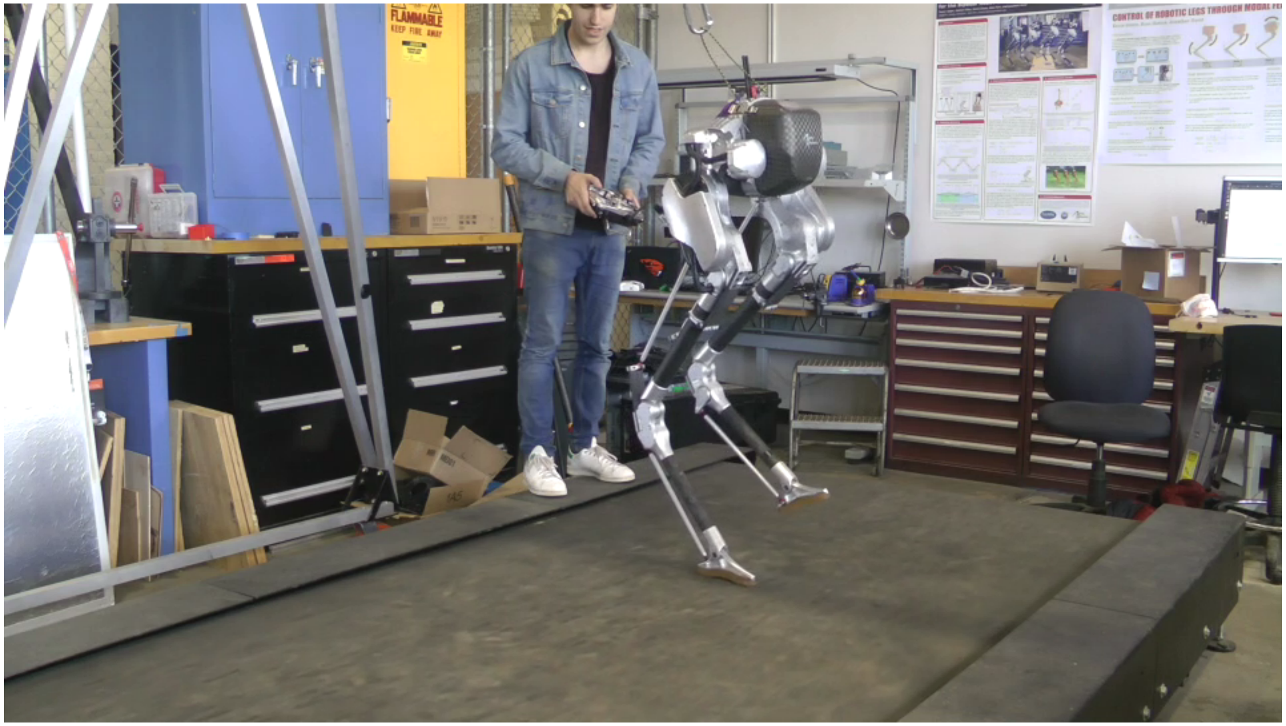

SMOOTH STYLE

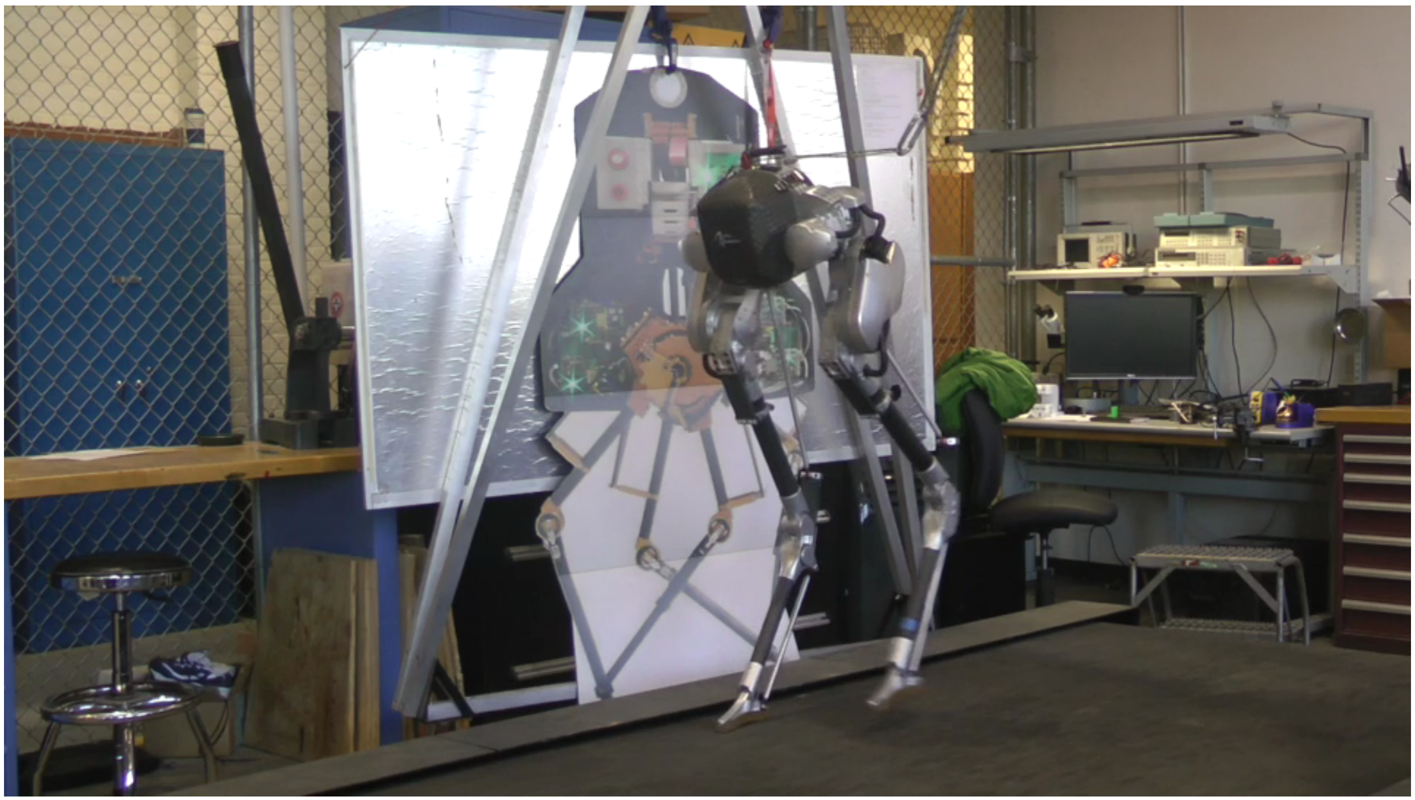

HIGH-STEPPING

PERTURBATION TEST

DEEP-MIMIC FOR BIOMECHANICAL MODELS

Scalable Muscle-Actuated Human Simulation and Control

SIGGRAPH 2019 Conditional Accept, Seoul National University

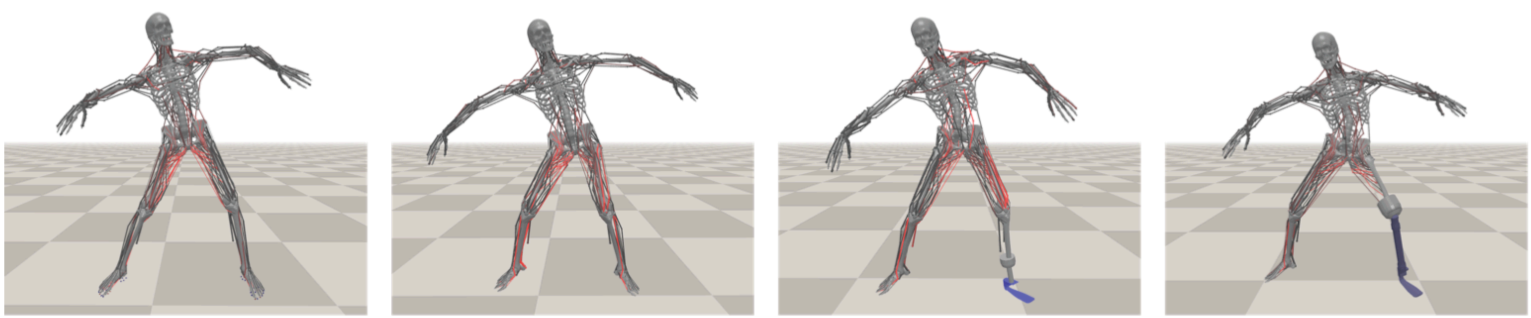

Scalable Muscle-Actuated Human Simulation and Control

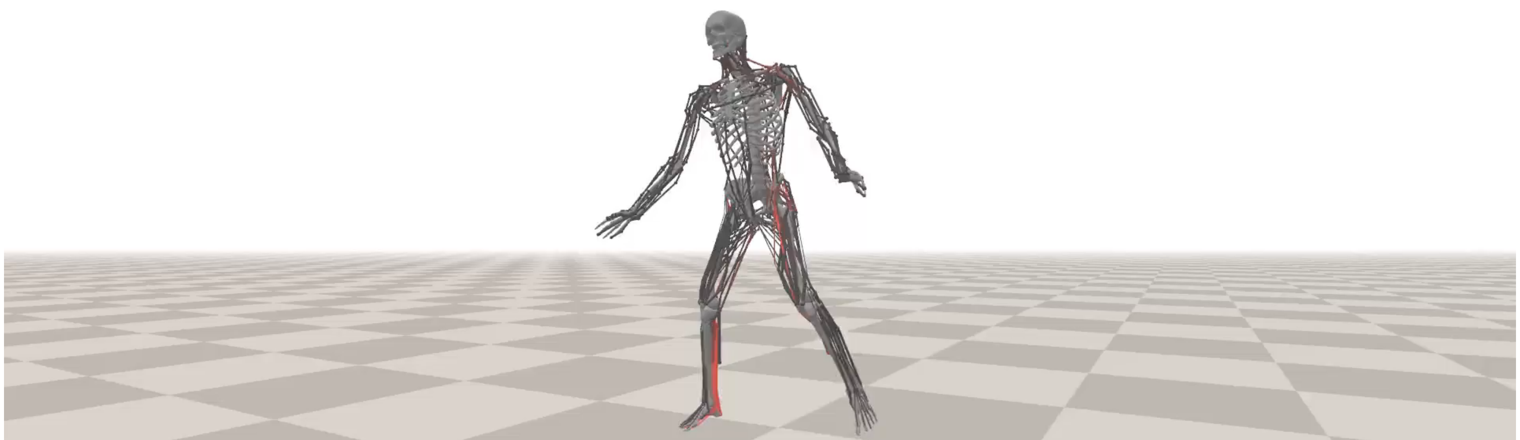

Seunghwan Lee(1), Kyoungmin Lee(2), Moonseok Park(2), and Jehee Lee(1)
Seoul National University(1), Seoul National University Bundang Hospital(2)

OVERVIEW

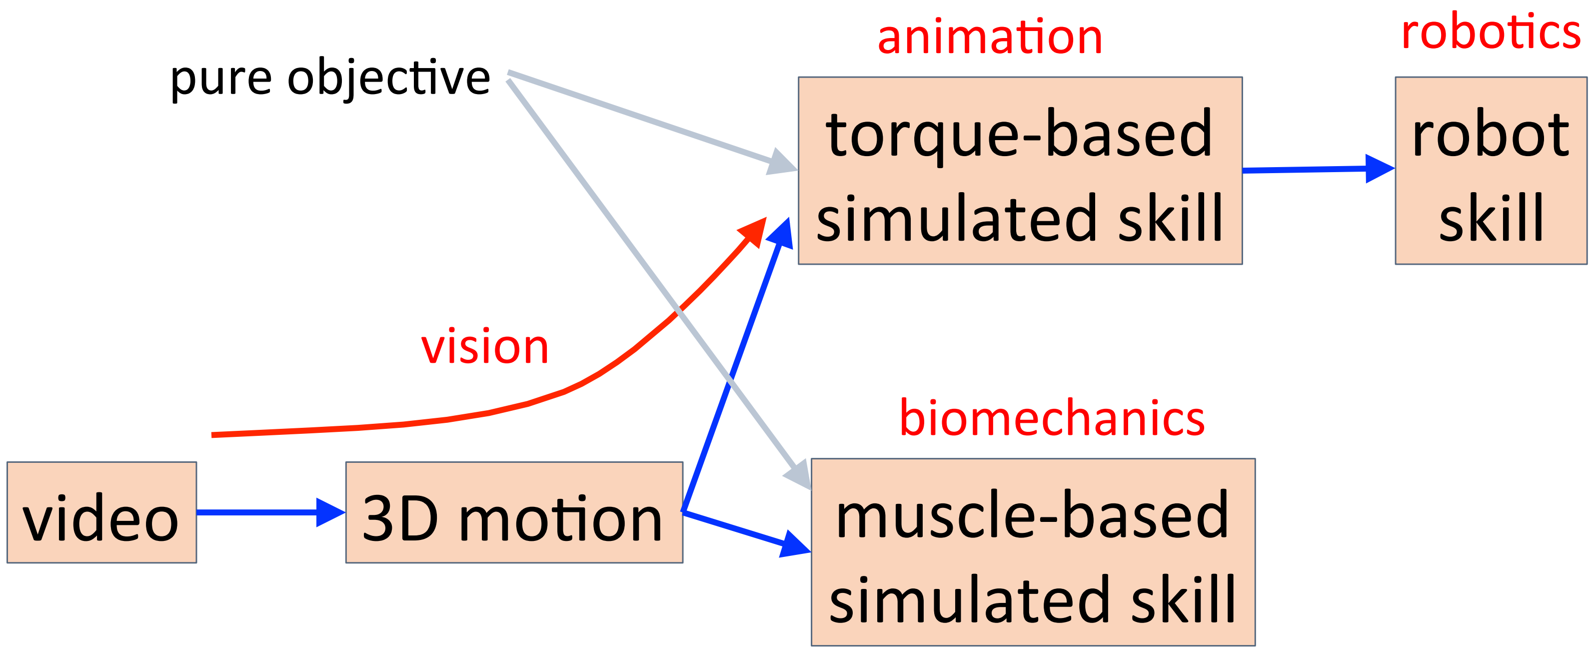
pure objective
animation
robotics
torque-based simulated skill
robot skill
vision
biomechanics
video
3D motion
muscle-based simulated skill

MUCH TO BE DONE…

- multi-skilled digital humans & animals
  - richer perception
  - interaction with the world
  - collaborative & interacting characters
- behavior from video
- motion planning

CONCLUSIONS

- “the best way to predict the future is to create it”
- physics-based human movement:
  rapid advances, many uses

ACKNOWLEDGEMENTS

Xue Bin Peng, Glen Berseth, Xinyi Zhang, Zhaoming Xie, Patrick Clary, Jonathan Hurst, Pieter Abbeel, Sergey Levine, Kangkang Yin, Hung Yu Ling

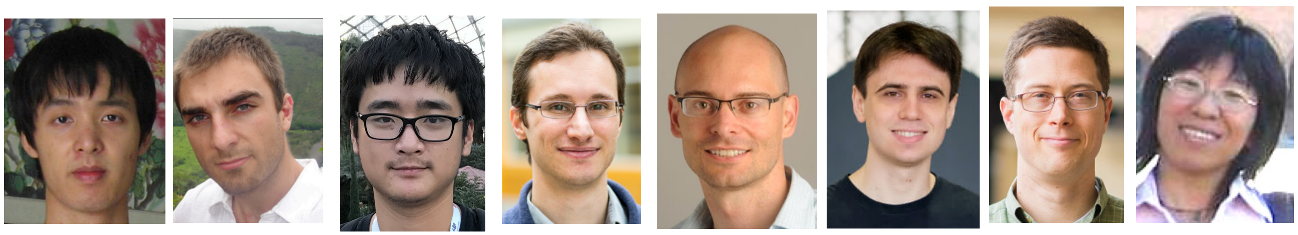

University of British Columbia

UC Berkeley

Oregon State University

Agility Robotics

Questions?